



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

(De-)Constructing TLS 1.3

Citation for published version:

Kohlweiss, M, Maurer, U, Onete, C, Tackmann, B & Venturi, D 2015, (De-)Constructing TLS 1.3. in *Progress in Cryptology - INDOCRYPT 2015 - 16th International Conference on Cryptology in India, Bangalore, India, December 6-9, 2015, Proceedings*. Springer, pp. 85-102, 16th International Conference on Cryptology in India, Bangalore, India, 6/12/15. https://doi.org/10.1007/978-3-319-26617-6_5

Digital Object Identifier (DOI):

[10.1007/978-3-319-26617-6_5](https://doi.org/10.1007/978-3-319-26617-6_5)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Progress in Cryptology - INDOCRYPT 2015 - 16th International Conference on Cryptology in India, Bangalore, India, December 6-9, 2015, Proceedings

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



(De-)Constructing TLS

Markulf Kohlweiss¹, Ueli Maurer², Cristina Onete³, Björn Tackmann^{*,4}, and Daniele Venturi^{†5}

¹*Microsoft Cambridge*

²*ETH Zürich*

³*Inria/ IRISA Rennes*

⁴*UC San Diego*

⁵*Sapienza University of Rome*

April 22, 2015

Abstract

TLS is one of the most widely deployed cryptographic protocols on the Internet; it is used to protect the confidentiality and integrity of transmitted data in various client-server protocols. Its non-standard use of cryptographic primitives, however, makes it hard to formally assess its security. It is in fact difficult to use traditional (well-understood) security notions for the key-exchange (here: *handshake*) and the encryption/authentication (here: *record layer*) parts of the protocol due to the fact that, on the one hand, traditional game-based notions do not easily support composition, and on the other hand, all TLS versions up to and including 1.2 combine the two phases in a non-standard way.

In this paper, we provide a modular security analysis of the handshake in TLS version 1.2 and a *slightly sanitized version* of the handshake in the current draft of TLS version 1.3, following the constructive cryptography approach of Maurer and Renner (ICS 2011). We provide a deconstruction of the handshake into modular sub-protocols and a security proof for each such sub-protocol. We also show how these results can be combined with analyses of the respective record layer protocols, and the overall result is that in all cases the protocol constructs (unilaterally) secure channels between the two parties from insecure channels and a public-key infrastructure. This approach ensures that (1) each sub-protocol is proven in isolation and independently of the other sub-protocols, (2) the overall security statement proven can easily be used in higher-level protocols, and (3) TLS can be used in *any* composition with other secure protocols.

In more detail, for the key-exchange step of TLS 1.2, we analyze the RSA-based and both Diffie-Hellman-based variants (with static and ephemeral server key share) under a non-randomizability assumption for RSA-PKCS and the Gap Diffie-Hellman assumption, respectively; in all cases we make use of random oracles. For the respective step of TLS 1.3, we prove security under the Decisional Diffie-Hellman assumption in the standard model. In all statements, we require additional standard computational assumptions on other primitives. In general, since the design of TLS is not modular, the constructive decomposition is less fine-grained than one might wish to have and than it is for a modular design. This paper therefore also suggests new insights into the intrinsic problems incurred by a non-modular protocol design such as that of TLS.

^{*}Part of the work done while at ETH Zürich. Author is supported by the Swiss National Science Foundation (SNF).

[†]Part of the work done while at Aarhus University supported by the Danish Council for Independent Research via DFF Starting Grant 10-081612.

Contents

1	Introduction	3
1.1	Overview and Previous Work	3
1.2	Contributions of this Paper	5
1.3	Outline	7
2	Preliminaries and Notation	11
2.1	Notation	11
2.2	Constructive Cryptography	11
2.3	Abstract Systems	11
2.4	The Notion of Construction	13
2.5	Discrete Systems	15
2.6	Insecure Communication Channels and TLS Fragments	16
3	Constructing the Master Secret	16
3.1	The Assumed Resources	16
3.2	Session Naming	18
3.3	Constructing the Shared Key	21
4	Expanding the Key	45
4.1	Key Expansion in TLS 1.2	45
4.2	Key Expansion in TLS 1.3	50
5	Constructing a Unilaterally Secure Channel	52
5.1	Cipher Suites based on Stream Ciphers	53
5.2	Cipher Suites based on CBC Encryption	54
5.3	Cipher Suites based on AEAD Encryption	56
6	Reconstructing TLS	57
6.1	Full Security Statements	59
7	Conclusion and Lessons Learned	61
A	More Details on TLS	66
A.1	X.509 Certificates	66
A.2	RSA PKCS#7	67
A.3	Key Expansion	68
B	Further Notation and Preliminaries	68
B.1	Signature Schemes	69
C	Game-based Definitions	70
C.1	OW-PCA & NR-PCA	70
C.2	Gap Diffie-Hellman	71
C.3	Collision Resistance	72
C.4	Unforgeability under Chosen-Message Attacks	72
C.5	Authenticated Encryption with Associated Data	72

1 Introduction

One of the most important applications of cryptography is the establishment of secure communication channels between two entities (e.g. a client and a server), and the protocol most widely used for this purpose is SSL/TLS. A key goal of research in cryptography is to provide security proofs for cryptographic protocols. This task is particularly difficult if the considered protocol has not been designed with provable security in mind, as is the case for SSL/TLS.

Initially developed by Netscape as the Secure Socket Layer (SSL) protocol [Hic95], the SSL/TLS protocol family suffered from several vulnerabilities; this led to the development of a number of subsequent protocol versions, each one fixing flaws discovered in the previous version. The most recent standardized protocol version is known as Transport Layer Security (TLS) version 1.2 [DR08], but there is also a working draft for version 1.3 [DR15] which modifies the protocol considerably. While the initial protocol was developed for protecting HTTP connections between a browser and a web server, many current Internet protocols including, e.g., SMTP or IMAP for transmitting e-mails and LDAP for accessing directories, have been extended to allow for securing the transmission with TLS. Because of its practical importance, the security of SSL/TLS has attracted a lot of attention in the literature, see e.g. [WS96, Pau99, Kra01, CK02, JK02, Bar04, HSD⁺05, MSW08, MT10, PRS11, AP12, BFCZ12, FHM⁺12, GIJ⁺12, JKSS12, ABP⁺13, AP13, BFK⁺13a, BFS⁺13, GKS13, KSS13, KPW13, BMM⁺15], in chronological order.

1.1 Overview and Previous Work

The TLS protocol consists of two parts, the *handshake*—essentially a key-exchange protocol that can be used with either unilateral or mutual authentication—and the *record protocol*—which protects the transmission of application data using the key obtained during the handshake. The TLS handshake offers several alternative key-exchange methods based on different cryptographic primitives. In each session, the actual method is chosen depending on the implementation, the available public keys, and the configuration. In all TLS versions up to and including 1.2, there are three standard methods: based on RSA encryption, on a static certified Diffie-Hellman key, or on signatures and an (ephemeral) Diffie-Hellman exchange, respectively. The draft for TLS 1.3 does away with RSA and static Diffie-Hellman, and uses a sanitized version of the ephemeral Diffie-Hellman based protocol. The typical setting is to only authenticate the server to the client, but client authentication is also possible if the client has a certified public key.

While a sequence of results about the record protocol [Kra01, MT10, PRS11, BMM⁺15] provides a comprehensive treatment thereof, the handshake protocol is not as well-understood. The reason is that the protocol was not designed with provable security in mind (see details in Section 1.2); it is inherently non-modular and uses cryptographic primitives in non-standard ways, which severely complicates security proofs. One particular observation with respect to non-modularity appearing in many analyses is that, in all versions up to 1.2, the final messages of the handshake already *use* the keys that are agreed upon;¹ this fact that prohibits an analysis of the full handshake in common security models (e.g., [BR93, CK01, CK02]). Our analysis of the initial part of the handshake can be interpreted as that the handshake without the final message can indeed be seen as a key-exchange protocol. Comparing this definition to previous models, our analysis covers the (game-based) notion of “key revealing” (because the distinguisher obtains all generated keys), but no “adaptive state reveals” defined in some game-based models.²

Beyond that, because of the non-standard use of schemes and primitives, papers that analyze

¹In the draft of TLS 1.3, specific keys are used for the protection of the final message.

²TLS-RSA and TLS-DH are in fact insecure with respect to state reveals.

(parts of) TLS 1.2 must generally choose between analyzing a modified version of the protocol, analyzing the original protocol in idealized models (such as the random oracle model), or using tailor-made computational assumptions.

Proving modified protocols was an early approach towards obtaining an intuition of the security of the TLS handshake. For example, [JK02] drop the final message of the handshake, while [MSW08] do consider the final messages, but in an unencrypted form. A more extreme approach is taken by Gajek *et al.* [GMP⁺08], who modify several details to achieve security in the UC framework [Can01]. (As pointed out by Küsters *et al.* [KT11] this additionally requires the insertion of session identifiers into the protocol. Küsters *et al.* also give a case study and describe how one could model TLS more faithfully in the IITM [KT11] model.) However, extending the security proof of a *modified* protocol to that of the *real* protocol is usually hard; for SSL/TLS, it is nearly impossible.

Recent game-based analyses of TLS made important steps towards assessing the security of the actual protocol. Jager *et al.* [JKSS12] treated the security of a specific configuration for the key-exchange step, i.e. TLS-DHE. In their (game-based) analysis, Jager *et al.* consider both the TLS handshake and the record protocol to be executed together. By treating the last message of the handshake as the first message of the record protocol, they circumvent the problems related to using one of the extracted keys for encryption during the key-exchange protocol. They showed that the TLS handshake and the record protocol *together* achieve the tailor-made (game-based) security notion “authenticated and confidential channel establishment,” under a specific, new hardness assumption. A (slightly) more modular analysis is given by Brzuska *et al.* [BFS⁺13], who view the TLS handshake and record protocol as a composition of a key agreement protocol and a scheme for achieving secure communication. They weaken the usual key-exchange security definitions and show explicitly that, under a specific notion of composition, the two parts can be proved secure together, in a meaningful way. Thus, this approach enables a security proof with some degree of modularity.

The recent work of Krawczyk, Paterson, and Wee [KPW13] extends the analysis of TLS-DHE [JKSS12] to other handshake configurations, like TLS-DH and TLS-RSA, and it is to our knowledge the most comprehensive analysis of TLS to date. Their approach is based on modeling the TLS handshake by means of a Key Encapsulation Mechanism (KEM). They also describe how to instantiate the KEM to capture each configuration. Subsequently, they show that, if the KEM attains a security notion known as Constrained CCA security (CCCA), the combined protocol consisting of KEM and record layer attains the ACCE security notion introduced by Jager *et al.* [JKSS12].

The same approach of using key encapsulation and key derivation was also taken recently by Bhargavan *et al.* [BFK⁺13b]. In a work parallel to ours, they use the notion of agility and game-based definitions in order to analyze a realistic form of TLS, including ciphersuite negotiation, renegotiation, and resumption. This work co-evolved with ours; while their analysis is closer to the implementation of the protocol, the advantage of our approach lies in the modularity of our deconstruction and proofs. Additionally, our definitions also apply to TLS 1.3, which does not hold for the KEM-based analyses of [KPW13, BFK⁺13b].

Results on using tools to verify TLS. Other approaches have used tools from mechanized logic to analyze TLS. For example, the symbolic model approach of [Pau99, HSD⁺05, BFCZ12]. [BFCZ12] already experimented with an executable specification for TLS. In addition to the purely symbolic PROVERIF, they started to use CRYPTOVERIF, a computationally faithful tool, but for a rather coarse security model. Recently, [BFK⁺13a] broke with symbolic models and gave a meaningful formal statement about mTLS, a standard compliant implementation of TLS. They describe cryptographic idealizations of the cryptographic primitives of TLS which

are indistinguishable by typed adversaries. This allows for a type-based and cryptographically faithful verification of mTLS using a dialect of F^* [SCF⁺11]. [BFK⁺13b] builds on [BFK⁺13a] and focuses on the details of the TLS handshake using the EASYCRYPT tool in addition to F^* .

As both EASYCRYPT and constructive cryptography build on a simple, modular semantic theory, we are planning on investigating the formal relation between constructive cryptography and EASYCRYPT’s new module system.

Previous results in constructive cryptography. Some related schemes have already been analyzed following the constructive cryptography paradigm [MR11, Mau11]. Maurer and Tackmann [MT10] used this approach to analyze the TLS record protocol as a construction of a secure channel from a shared key and insecure communication channels. Recently, Maurer et al. [MTC13] described a unilateral key-exchange protocol in this model; the goal of their protocol is similar to the goal of the TLS configurations without client authentication; their protocol is considerably simpler and secure under weaker assumptions, but does not allow to adaptively choose cipher suites. The new key-exchange mode in TLS 1.3 follows the same ideas as the protocol analyzed in [MTC13], which is based on a protocol called A-DHKE by Shoup [Sho99].

1.2 Contributions of this Paper

In this paper, we define and prove the security of TLS version 1.2 and a *slightly sanitized version* of the drafted version 1.3 following the constructive cryptography paradigm of Maurer and Renner [MR11]. To the best of our knowledge, we are the first ones to provide a provable-security treatment of the drafted protocol TLS 1.3—with the caveat that we modify the protocol to *not* encrypt the server’s certificate. We note that this change seems necessary for any kind of security analysis, since encrypting the certificate (which is needed in the key computation) introduces a circularity in the proof. We also note that the absence of the encryption does not diminish the *security* of the resulting channel, though it may decrease the *privacy* guarantees *for the server*. In the case of unilateral authentication, however, the maximum privacy that can be achieved by TLS is anyway only against a passive attacker.

In constructive cryptography, the (security) guarantees provided to parties in a specific context are formalized in terms of *resources* available to the parties, typically channels with certain properties. Cryptographic protocols *construct* (desired) resources from assumed resources, and the composition theorem of this framework guarantees that the protocol (using the resources assumed by it) can then be used whenever the constructed resource is required (as an assumed resource) in future constructions. This approach can be viewed as being akin to the OSI communication model, in which the protocol(s) used at every OSI layer are handled in isolation, and are then incorporated into the next layer of processing using a well-defined interface between the two. In the same way, in constructive cryptography each construction is handled in isolation, and the constructed resource can then be used in further constructions. Since the actual security statement for the overall protocol is of a standardized form, in terms of a resource, it is straightforward to use the protocol in a higher-level context, with the overall security proof again following from the composition theorem.

This methodology favors a modular approach to provable security, in which even a complex protocol can be de-constructed in layers, each of which can be analyzed in isolation, finally using the composition theorem to prove the soundness of the entire approach. We show its power in de-constructing the TLS/SSL protocol, and showing that our approach can easily be used to prove that both TLS 1.2 (in all three modes of operation), and a slightly modified version of TLS 1.3 construct unilaterally secure channels between clients and a server.

In this spirit, we show that SSL/TLS with unilateral authentication constructs a *unilaterally*

secure communication channel from an insecure communication network such as the Internet, a public-key infrastructure, and—in the case of TLS 1.2—a public random oracle (an idealization of a hash function). We de-compose the TLS protocol into a sequence of sub-protocols that each achieve a smaller construction step. Each step is then achieved by one or more cryptographic schemes (like key-exchange mechanisms, key confirmation, symmetric encryption, and MAC), or by security mechanisms (like nonces). The goal of each scheme or mechanism is to *construct* an ideal resource, which then serves as an assumed resource in the next step. This approach can be seen as an analogue to the layered structure in the network protocol stack, where for instance IP provides end-to-end connectivity—constructing pairwise unreliable channels—and TCP constructs reliable channels from these unreliable channels, and an application protocol like FTP constructs a resource like file storage from the reliable communication channels.

The composition of all steps we analyze yields the complete TLS protocol, and the composition theorem guarantees the soundness of this approach. This method of analyzing the TLS protocol has two main advantages:

1. The composition theorem (cf. [MR11, MT10, KMO⁺13]) guarantees that the TLS protocol can be used, given the assumed resources, whenever an application or higher-level protocol requires the type of secure channel described in our definitions.
2. As each step is proven in isolation and the steps are combined generically by the composition theorem,
 - (a) the analysis of each individual step becomes simpler, because it can be performed independently of the other steps;
 - (b) if an alternative security analysis of one step is given, then one can immediately use this proof without adapting the proofs of the remaining primitives. In particular, by adapting the techniques of [JKSS12, KPW13] to prove the key-exchange steps without a random oracle, we can immediately obtain a standard-model analysis of the complete protocol without re-proving the remaining steps;
 - (c) further optional sub-protocols, like other symmetric ciphers or elliptic curve methods, are integrated into our analysis by proving simply that they construct a specific step. The security of the full protocol will then follow generically.
 - (d) Construction steps that are proven in isolation can directly be used in other protocols. For example, our statement that the complete TLS protocol constructs a unilaterally secure channel can be combined with the approach of authenticating the client by transmitting a password as described in [MTC13].

Unfortunately, TLS tries hard to resist this modularization, and it also uses many techniques heuristically: the cryptographic keys intended to protect the payload are also used for the confirmation (here called “finished”) messages, the pseudo-random function is keyed with a non-uniform key, and various protocol messages are unnecessarily included in the confirmation messages, to name a few. These unfortunate design choices complicate our analysis: sometimes abstractions cannot be introduced at the intuitively appealing levels, often “auxiliary” data must be passed through multiple protocol layers, and as the potential interference of different protocol sessions is resolved only at a late stage, many protocol steps must be analyzed in a complicated “multi-session” scenario.

Furthermore, it appears impossible to perform a reasonably modular analysis for the RSA-PKCS variant of the handshake under the OW-PCA assumption [JK02], even in the random oracle model. This is evidence that it is difficult to modularize the proof of the CCA security of the TLS key extraction mechanism of [KPW13] further unless one is willing to make additional assumptions. Building on a recent result by [BFK⁺13b], we show that a reasonably weak non-randomizability assumption that they call NR-PCA is sufficient for a more modular constructive

proof. We describe these problems in more detail in Section 3.3.3.

Although various of the above mentioned obstacles could have been bypassed by modifying the scheme, we kept our analysis close to the realistic TLS protocol. The result is a more complicated analysis; however, the artifices are immanent to the protocol, not to our technique.

In total, we analyze the three main methods for the establishment of the master secret in the handshake of TLS 1.2, namely TLS-DH, TLS-DHE, and TLS-RSA, as well the two main methods for protecting the payload messages in the record layer protocol, namely the Authenticate-then-Encrypt combinations using a stream cipher resp. a block cipher in CBC-mode, and a MAC. By the composition theorem, we obtain the security of all possible combinations. We use the GapDH assumption for the Diffie-Hellman based cipher suites, the NR-PCA assumption for RSA PKCS#7, pseudo-randomness for the stream and block ciphers, and strong unforgeability for the MAC. Moreover, we require the hash function used in the handshake to be collision resistant, and our proofs of the handshake protocols are based on the random oracle assumption. For TLS 1.3, we use the Decisional Diffie-Hellman assumption in our standard-model proofs; we additionally require the hash function to be collision resistant, HMAC to be a PRF (with two different distributions for the keys), and the AEAD mode to be secure with respect to standard security notions.

On the type of proven statements. Results on provable security differ with respect to (1) the assumptions made and (2) the statement that is proved to follow from the assumptions. It is important that the proved statement is of a form that allows for both comparisons of protocol performance, and for direct use in the proof of a higher-level protocol. Security statements should thus be exact (as opposed to asymptotic), giving precise upper bounds for the security level guaranteed by a protocol. Furthermore, a key to analyzing and designing cryptographic protocols is a modularization in which the role of each cryptographic primitive (e.g. encryption) or mechanism (e.g. nonce exchange) is made explicit, and the security of its application is proved in isolation, once and for all. The constructive cryptography framework provides a sound instantiation of this approach.

As the systems model underlying our work has no notion of time, we currently cover neither timing attacks nor forward secrecy (sometimes called “key corruptions”). Note that TLS-RSA and TLS-DH are in fact insecure with respect to forward secrecy. Our security statements are concrete (rather than asymptotic). In particular, all statements are parametrized in the respectively relevant parameters such as the number of protocol sessions or the total length of the transmitted messages. We prove *exact* reductions to parametrized games which are used to formalize computational hardness of the underlying primitives; given concrete assumptions on the computational hardness, our statements imply concrete bounds for the security of the TLS protocol.

1.3 Outline

We now describe, at a high level, the cryptographic mechanisms used in TLS and the constructive steps they achieve. We start with a simplified description of the protocol in Figure 1, where the notation $\langle m \rangle_\kappa$ denotes the encryption of a message m , possibly together with other messages, under the key κ .

Decomposition into sub-protocols. At a high level, our analysis proceeds by repeatedly “scraping off” a part of the protocol, starting at the beginning; at each step, we consider the remainder of the protocol as the “payload” of the part we analyze. This permits us to then

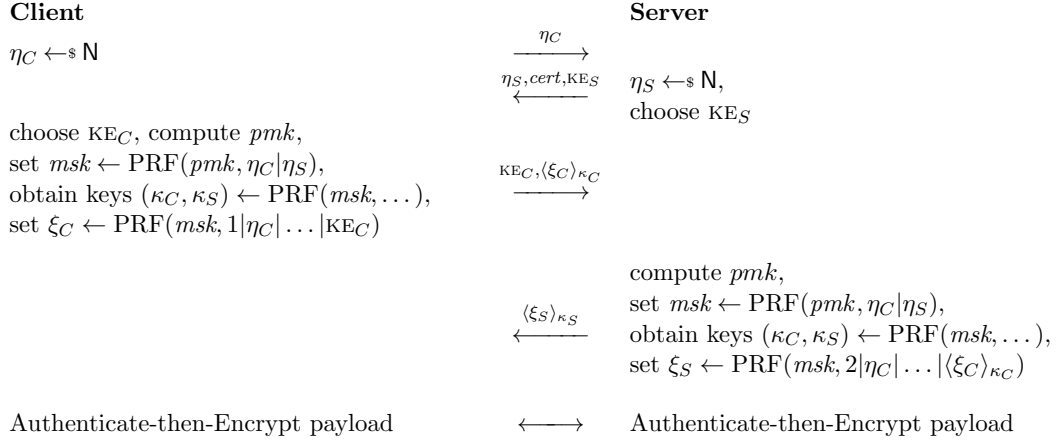


Figure 1: The TLS 1.2 handshake as a message exchange

use the composition theorem to conclude the security of the entire protocol. We proceed in the following steps.

1. *Nonce generation.* Each client chooses a nonce $\eta_C \leftarrow \text{\$ } \mathbf{N}$ for a distribution \mathbf{N} on the set of nonces.

Since TLS is often used in a unilateral mode where a client has no distinguished name, the client nonce can be seen as a “disposable name” that is (with high probability) unique to that client—the purpose of this nonce is to separate protocol sessions of different clients.

This guarantee is formalized as a resource **NAME**, which provides a unique name to each client.

2. *Nonce exchange.* The client sends the nonce η_C (which it obtained from the resource **NAME**). The server chooses a nonce $\eta_S \leftarrow \text{\$ } \mathbf{N}$ and sends it to the client.

The server uses the client nonce to identify the client—note that we focus on the unilateral case where the client is not authenticated. The server nonce is less crucial in our analysis, but to guarantee in the key derivation step (for which we make a random oracle assumption) that the data in each session is unique, we assume uniqueness of the server nonce, losing a collision term. However, if one analyzes the key derivation in more detail and not in the random oracle model, the entropy contained in the server nonce may be a valuable input for extraction.

This guarantee is formalized as a resource **SNET** that outputs nonces (potentially chosen adversarially) and associates them with the (fully) insecure communication channels giving no guarantee of consistency.

3. *Unilateral key exchange.* The server chooses a key-exchange message KE_S according to the chosen cipher suite (the message may be empty or a Diffie-Hellman element), and sends its certificate cert and KE_S to the client. The client also chooses a key-exchange message KE_C accordingly and sends it to the server. Both compute the pre-master key pmk according to the chosen scheme, and compute the master secret key msk (at this step, we use a random oracle instead of a PRF).

The core key-exchange protocol results in a pre-master secret, but the employed mechanisms are too weak to achieve a meaningful security guarantee for this step alone. Obtaining the master secret key msk by querying the random oracle “de-correlates” the keys obtained in different sessions and also normalizes their distribution.

The complete step constructs a resource **MSK** that outputs uniform random keys (the msk) to the client and the server; the keys may, however, be different in case the adversary interfered.

4. *Key expansion and generating confirmations.* Both client and server compute keys κ_C and κ_S and “finished” messages ξ_C and ξ_S via a pseudo-random function PRF keyed with msk . The two finished messages are then encrypted, in a concatenation of all the messages in the protocol, and sent.

Since the transmission of ξ_C and ξ_S in the protocol is protected by Authenticate-then-Encrypt with the same keys as the payload messages, a modularization at the intuitively appealing point *after* performing the confirmation is impossible. Only after this protocol step can we fully separate different “sessions,” as the computation of ξ_C and ξ_S still requires data which is correlated between different sessions (e.g., the server certificate).

The behavior of this resource is similar to that of **MSK** with a different key space, but the resource can now be described as the parallel composition of one (single-session) resource $\stackrel{KSP,*}{=} \bullet$ per client.

5. *Channel setup and payload transmission.* The channel setup consists of sending and verifying the messages ξ_C and ξ_S protected with Authenticate-then-Encrypt with keys κ_C and κ_S , respectively. If the verification succeeds, the client and server exchange payload messages protected by Authenticate-then-Encrypt.

The resulting resource $\leftarrow^* \rightarrow \bullet$ is, for each client, one channel to the server which (at the decision of the adversary) allows for either fully secure client-server communication, or for server-adversary communication. The client notices which case occurs, the server does not.

Constructions achieved by the sub-protocols. We sketch the decomposition in Figure 2. The resources **NET** (the network), **PKI** (the PKI), and **RO** (the random oracle) used by the protocol are depicted on the right hand side, and the “converters” (each capturing a different part of the protocol) are drawn as boxes. The first type of converter **rnd** generates a random nonce. The sub-protocol consisting of one such converter per client constructs the resource **NAME** that outputs a *unique* nonce to each client (modulo a collision term, which becomes explicit in the security statement).

The second converter **hec** uses as resources **NAME** and the network **NET**, exchanges nonces with the server via **NET**, and also transmits “payload” messages from higher-level protocols. Together with the server’s converter **hes**, this constructs the “network with nonces,” denoted as **SNET**. The construction statement can be written as

$$[\mathbf{NET}, \mathbf{NAME}] \xRightarrow{\text{hec, hes}} \mathbf{SNET}.$$

The third converter is then alternatively one of **dhc** (for static DH), **dhec** (for ephemeral DH), and **rsac** (for RSA-PKCS) depending on the cipher suite used in the handshake. All these use the resources **PKI** and **RO**, as well as the resource **SNET** constructed by **rnd** and **(hec, hes)**, and construct, together with the server’s counterpart, a “master secret key” resource later called “**MSK**”.

The statements for the following layers are (here for the static DH scheme and a set \mathcal{C} of clients)

$$[\mathbf{SNET}, \mathbf{PKI}, \mathbf{RO}] \xRightarrow{\text{dhc, dhsg}} \mathbf{MSK}, \quad \mathbf{MSK} \xRightarrow{\text{expc, exps}} \bigotimes_{C \in \mathcal{C}} \left[\stackrel{KSP,*}{=} \bullet \right]^{(C, S/\eta_C)},$$

$$\text{and} \quad \stackrel{KSP,*}{=} \bullet \xRightarrow{\text{atec, ates}} \leftarrow^* \rightarrow \bullet.$$

In the above statements, $\stackrel{KSP,*}{=} \bullet$ is a resource which is specified with respect to a single client and the server, and the product together with the brackets signifies that there is an independent such

2 Preliminaries and Notation

2.1 Notation

For a number $n \in \mathbb{N}$, we use the notation $[n] := \{1, \dots, n\}$. Also, for a distribution X over a set \mathcal{X} , we write $X \leftarrow \$ \mathsf{X}$ to denote that the random variable X is sampled from distribution X , and we write $X \leftarrow \$ \mathcal{X}$ to express that X is sampled uniformly at random from the set \mathcal{X} . We write \mathcal{B} for the set $\{0, 1\}^8$ of bytes.

2.2 Constructive Cryptography

The foundational idea of constructive cryptography [MR11, Mau11] is to specify both the assumptions³ and the guarantees of protocols explicitly as resources, and to consider a protocol as a construction of a (desired) resource from assumed resources. A resource is a shared functionality accessed by several parties; in this work we consider different types of communication channels and shared keys. The assumed resources formalize the setting in which a protocol is used (such as a certain type of communication channel) and constructed resources describe the functionality achieved by using the protocol on the assumed resources (such as a shared key or a communication channel with stronger guarantees). A protocol in this setting is described as a tuple of so-called converters; one per (honest) party.

2.3 Abstract Systems

Resources and converters are modeled as systems. At the highest level of abstraction (following the hierarchy in [MR11]), *systems* are objects with *interfaces* by which they connect to (interfaces of) other systems; each interface is labeled with an element of some label set and connects to only a single other interface. Multiple interfaces can be merged into a single interface; the original interfaces are referred to as *sub-interfaces* of the composite interface. (Sub-interfaces have labels relative to the composite interface; to refer to a sub-interface S of some interface I , we write I/S .) This concept of *abstract systems* captures the topological structures that result when multiple systems are connected in this manner. It does not, however, model the behavior of systems, i.e., *how* the systems interact via their interfaces; statements about specific protocols are statements at the next (lower) abstraction level. In this work, we describe all systems in terms of (probabilistic) discrete systems.

Resources and converters. *Resources* in this work are systems with interfaces labeled by elements of some label set \mathcal{L} . A *converter* is a two-interface system which is directed in that it has an *inside* and an *outside* interface. As a notational convention, we generally use upper-case, bold-face letters (e.g., \mathbf{R} , \mathbf{S}), symbols (e.g., $\bullet \rightarrow$), or upper-case sans-serif fonts to denote resources, and lower-case Greek letters (e.g., α , β) or sans-serif fonts (e.g., `enc`, `dec`) for converters. We denote by $\Phi_{\mathcal{L}}$ (or simply Φ if \mathcal{L} is clear from the context) the set of all resources with interface labels in \mathcal{L} , and by Σ the set of all converters. We use two special converters, an “identity” converter id and a “blocking” converter \perp that has an inactive outside interface.

System composition. The topology of a composite system is described using a term algebra, where each expression starts from one resource on the right-hand side and is subsequently extended with further terms on the left-hand side. An expression is interpreted in the way that all interfaces of the system it describes can be connected to interfaces of systems which are

³The term “assumption” often refers to two different concepts: setup assumptions such as a network or a PKI, and computational assumptions such as the hardness of certain problems. Here, we refer to setup assumptions.

appended on the left. For instance, for a single resource $\mathbf{R} \in \Phi$, all its interfaces are accessible. For $I \in \mathcal{L}$, a resource $\mathbf{R} \in \Phi$, and a converter $\alpha \in \Sigma$, the expression $\alpha^I \mathbf{R}$ denotes the composite system obtained by connecting the inside interface of α to the I -interface of \mathbf{R} ; the outside interface of α becomes the I -interface of the composite system. The system $\alpha^I \mathbf{R}$ is again a resource. For two resources \mathbf{R} and \mathbf{S} , $[\mathbf{R}, \mathbf{S}]$ denotes the parallel composition of \mathbf{R} and \mathbf{S} . For each $I \in \mathcal{L}$, the I -interfaces of \mathbf{R} and \mathbf{S} are merged and become the *sub-interfaces* of the I -interface of $[\mathbf{R}, \mathbf{S}]$. A converter α that connects to the I -interface of $[\mathbf{R}, \mathbf{S}]$ has two inside sub-interfaces, where the first connects to \mathbf{R} and the second connects to \mathbf{S} (i.e., sub-interfaces are ordered). Finally, we denote by id a special converter that forwards all messages from the inside to the outside interface and vice versa, hence $\text{id}^I \mathbf{R} \equiv \mathbf{R}$.

We introduce special notation for families of resources or converters: If we compose a family of resources $(\mathbf{R}_i)_{i \in \{1, \dots, n\}}$ (resp. converters $(\alpha_i)_{i \in \{1, \dots, n\}}$) in parallel, we write this as a product such as $\bigotimes_{i=1}^n \mathbf{R}_i$ (resp. $\bigotimes_{i=1}^n \alpha_i$). If we attach a family of converters $\alpha_1, \dots, \alpha_n$ to interfaces I_1, \dots, I_n of a resource \mathbf{R} , we write $\prod_{i=1}^n \alpha_i^{I_i} \mathbf{R}$.

Each setting is described by a cryptographic algebra with a certain label set \mathcal{L} .

Definition 2. A *cryptographic algebra* for a label set \mathcal{L} is a pair $\langle \Phi, \Sigma \rangle$ consisting of a set of resources Φ and a set of converters Σ , together with families of parallel composition operations $[\cdot] : \Phi^* \rightarrow \Phi$ and $[\cdot] : \Sigma^* \rightarrow \Sigma$, as well as connecting operations $\cdot^i : \Sigma \times \Phi \rightarrow \Phi$. The algebra is *composition-order independent* if

1. for all $C_1, C_2 \in \Sigma$, $R \in \Phi$, and $i, j \in \mathcal{L}$ with $i \neq j$,

$$C_1^i(C_2^j R) = C_2^j(C_1^i R), \text{ and}$$

2. for all $C_1, \dots, C_m \in \Sigma$, $R_1, \dots, R_m \in \Phi$, and $i \in \mathcal{L}$,

$$[C_J]^i[R_J] = [C_1^i R_1, \dots, C_m^i R_m].$$

Distinguishers. A *distinguisher* \mathbf{D} is a special type of system that connects to all interfaces of a resource \mathbf{U} and outputs a single bit at the end of its interaction with \mathbf{U} . In the term algebra, this appears as the expression $\mathbf{D}\mathbf{U}$, which defines a binary random variable. The *distinguishing advantage of a distinguisher \mathbf{D} on two systems \mathbf{U} and \mathbf{V}* is defined as

$$\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) \doteq |\Pr(\mathbf{D}\mathbf{U} = 1) - \Pr(\mathbf{D}\mathbf{V} = 1)|.$$

The advantage of a class \mathcal{D} of distinguishers is defined as $\Delta^{\mathcal{D}}(\mathbf{U}, \mathbf{V}) \doteq \sup_{\mathbf{D} \in \mathcal{D}} \Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V})$. The distinguishing advantage measures how much the distribution of the output of \mathbf{D} differs when it is connected to either \mathbf{U} or \mathbf{V} . Intuitively, if no distinguisher (of a certain class) differentiates between \mathbf{U} and \mathbf{V} , they can be used interchangeably in any environment of that class (otherwise that specific environment could serve as a distinguisher).

Note that the distinguishing advantage is a pseudo-metric. In particular, it satisfies the triangle inequality, i.e., $\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{W}) \leq \Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) + \Delta^{\mathbf{D}}(\mathbf{V}, \mathbf{W})$ for all resources \mathbf{U} , \mathbf{V} , and \mathbf{W} and distinguishers \mathbf{D} . There is an *equivalence* relation on the set of resources (which is defined on the level of discrete systems), denoted by $\mathbf{U} \equiv \mathbf{V}$, which means that $\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) = 0$ for *all* distinguishers \mathbf{D} .

Games. Games that capture properties such as unforgeability are two-interface systems that at their left interface connect to some adversary or solver \mathbf{A} and at the right interface output a single bit (usually denoted W). The performance of \mathbf{A} in a game \mathbf{G} is denoted as

$$\Gamma^{\mathbf{A}}(\mathbf{G}) \doteq \Pr^{\mathbf{A}\mathbf{G}}(W = 1).$$

Reductions. When relating (the hardness of) two problems such as distinguishing systems or winning a game, it is convenient to use a special type of system \mathbf{C} that translates one setting into the other. Formally, \mathbf{C} is a converter that has an *inside* and an *outside* interface. When it is connected to a system \mathbf{S} , which is denoted by \mathbf{CS} , the inside interface of \mathbf{C} connects to the merged interfaces of \mathbf{S} and the outside interface of \mathbf{C} becomes the interface of the composed system. \mathbf{C} is called a *reduction system* (or simply *reduction*).

To reduce distinguishing two systems \mathbf{S}, \mathbf{T} to distinguishing two systems \mathbf{U}, \mathbf{V} , one describes a reduction \mathbf{C} such that $\mathbf{CS} \equiv \mathbf{U}$ and $\mathbf{CT} \equiv \mathbf{V}$. Then, for all distinguishers \mathbf{D} , we have $\Delta^{\mathbf{D}}(\mathbf{U}, \mathbf{V}) = \Delta^{\mathbf{D}}(\mathbf{CS}, \mathbf{CT}) = \Delta^{\mathbf{DC}}(\mathbf{S}, \mathbf{T})$. The last equality follows from the fact that \mathbf{C} can also be thought of as being part of the distinguisher.

2.4 The Notion of Construction

The *construction* notion (cf. [Mau11]) requires two conditions: *availability* and *security*. For the former, the behavior of the real resource (with the converters) must be indistinguishable to the behavior of the ideal resource, if the “blocking” converter \perp is attached at the adversarial interface (always denoted by the special label E) of both resources. For the latter requirement, we demand the existence of a simulator which essentially emulates the behavior at the E -interface of the real resource, while being connected to the corresponding interface of the ideal resource. More formally:

Definition 3. Let $\Phi_{\mathcal{L}}$ and Σ be as above, and let ε_1 and ε_2 two functions mapping each distinguisher \mathbf{D} to a real number in $[0, 1]$. A protocol $\pi_{\mathcal{L}'} = (\pi_{\ell})_{\ell \in \mathcal{L}'}$ *constructs resource* $\mathbf{S} \in \Phi$ *from resource* $\mathbf{R} \in \Phi$ *with distance* $(\varepsilon_1, \varepsilon_2)$ *and with respect the simulator* σ , denoted

$$\mathbf{R} \xrightarrow[\Longleftrightarrow]{\pi_{\mathcal{L}'}, \sigma, (\varepsilon_1, \varepsilon_2)} \mathbf{S},$$

if, for all distinguishers \mathbf{D} ,

$$\begin{cases} \Delta^{\mathbf{D}}\left((\pi_{\mathcal{L}'})^{\mathcal{L}'} \pi_S^S \perp^E \mathbf{R}, \perp^E \mathbf{S}\right) & \leq \varepsilon_1(\mathbf{D}) & \text{(availability),} \\ \Delta^{\mathbf{D}}\left((\pi_{\mathcal{L}'})^{\mathcal{L}'} \pi_S^S \mathbf{R}, \sigma^E \mathbf{S}\right) & \leq \varepsilon_2(\mathbf{D}) & \text{(security),} \end{cases}$$

where $(\pi_{\mathcal{L}'})^{\mathcal{L}'}$ means attaching each π_{ℓ} to interface ℓ .

In the descriptions of the resources, we consider two different “modes:” one where the converter \perp is attached to the E -interface—this means that *no* attacker is present and formalizes the availability/correctness of the protocol—and one where the converter is not attached—this means that an attacker *is* present and may be attacking the communication; this condition formalizes the security of the protocol.

The composition theorem. The statements we prove each show the security of one protocol step in isolation, i.e. they prove for one protocol that it constructs some resource from one or more assumed resources, possibly under some (computational) assumption. The composition theorem now states that two such construction steps can be composed: if one (lower-level) protocol constructs the resource that is assumed by the other (higher-level) protocol, then the composition of those two protocols constructs the same resource as the higher-level protocol, but from the resources assumed by the lower-level protocol, under the assumptions that occur in (at least) one of the individual security statements.

To state the theorem, we extend the notation for parallel and sequential composition to protocols, i.e., we write $\psi_{\mathcal{L}'} \circ \pi_{\mathcal{L}'}$ or $[\pi_{\mathcal{L}'}^{(1)}, \dots, \pi_{\mathcal{L}'}^{(m)}]$ and mean that the respective operations

apply to all converters individually. We also make use of the special converter id that behaves transparently (i.e., allows access to the underlying interface of the resource). The protocol where all parties have to converter id is denoted $\text{id}_{\mathcal{L}'}$.

The composition theorem was first explicitly stated in [MT10], but the statement there was restricted to asymptotic settings. Later, in [KMO⁺13], the theorem was stated in a way that also allows to capture concrete security statements. The proof, however, still follows the same steps as the one in [MT10]. For the statement of the theorem we assume the operation $[\cdot, \dots, \cdot]$ to be left-associative; in this way we can simply express multiple resources using the single variable \mathbf{U} .

Theorem 4. *Let $\mathbf{R}, \mathbf{S}, \mathbf{T}, \mathbf{U} \in \Phi_{\mathcal{L}}$ be resources, and let $\mathcal{L}' = \mathcal{L} \setminus \{E\}$. Let $\pi_{\mathcal{L}'}$ and $\psi_{\mathcal{L}'}$ be protocols, σ_{π} and σ_{ψ} be simulators, and $(\varepsilon_{\pi}^1, \varepsilon_{\pi}^2)$, $(\varepsilon_{\psi}^1, \varepsilon_{\psi}^2)$ such that*

$$\mathbf{R} \xrightarrow{\pi_{\mathcal{L}'}, \sigma_{\pi}, (\varepsilon_{\pi}^1, \varepsilon_{\pi}^2)} \mathbf{S} \quad \text{and} \quad \mathbf{S} \xrightarrow{\psi_{\mathcal{L}'}, \sigma_{\psi}, (\varepsilon_{\psi}^1, \varepsilon_{\psi}^2)} \mathbf{T}.$$

Then

$$\mathbf{R} \xrightarrow{\alpha_{\mathcal{L}'}, \sigma_{\alpha}, (\varepsilon_{\alpha}^1, \varepsilon_{\alpha}^2)} \mathbf{T}$$

with $\alpha_{\mathcal{L}'} = \psi_{\mathcal{L}'} \circ \pi_{\mathcal{L}'}$ by composing all converters, $\sigma_{\alpha} = \sigma_{\pi} \circ \sigma_{\psi}$, $\varepsilon_{\alpha}^1(\mathbf{D}) = \varepsilon_{\pi}^1(\mathbf{D}\psi_{\mathcal{L}'}) + \varepsilon_{\psi}^1(\mathbf{D})$, and $\varepsilon_{\alpha}^2(\mathbf{D}) = \varepsilon_{\pi}^2(\mathbf{D}\psi_{\mathcal{L}'}) + \varepsilon_{\psi}^2(\mathbf{D}\sigma_{\pi}^E)$, where $\mathbf{D}\psi_{\mathcal{L}'}$ and $\mathbf{D}\sigma_{\pi}^E$ mean that \mathbf{D} applies the converters at the respective interfaces. Moreover

$$[\mathbf{R}, \mathbf{U}] \xrightarrow{[\pi_{\mathcal{L}'}, \text{id}_{\mathcal{L}'}], [\sigma_{\pi}, \text{id}], (\varepsilon_{\pi}^1, \varepsilon_{\pi}^2)} [\mathbf{S}, \mathbf{U}],$$

with $\varepsilon_{\pi}^i(\mathbf{D}) = \varepsilon_{\pi}^i(\mathbf{D}[\cdot, \mathbf{U}])$, where $\mathbf{D}[\cdot, \mathbf{U}]$ means that the distinguisher emulates \mathbf{U} in parallel.

Settings and interface sets considered in this work. The most important scenario we consider in this work comprises multiple clients, one server, and one (explicit, external) adversary. Some resources are easier described as the parallel composition of resources for fewer parties and some protocol steps can be proven in such a simpler setting in isolation, i.e., with respect to only one client, one server, and the adversary.

The simplified setting involving only two honest parties and one attacker is called the $\{A, B, E\}$ -setting and is used to analyze protocols and schemes like symmetric encryption or MACs, cf. [MT10]. The (honest) parties' interfaces are named A and B , and there is an explicit adversarial interface E . Resources are in the set $\Phi_{\{A, B, E\}}$, and protocols are pairs of converters $\pi = (\pi_1, \pi_2)$ for A and B , respectively.

Unilateral key-exchange protocols are used in a setting with multiple clients, one server, and an explicit adversary. We consider a set \mathcal{C} of clients, a server S , and an adversary E . Hence, we consider a label set $\mathcal{L} = \mathcal{C} \cup \{E, S\}$, resources are in the set $\Phi_{\mathcal{L}}$, and a protocol consists of a family $(\pi_C)_{C \in \mathcal{C}}$ of client converters and a server converter π_S .

Constructions in the $\{A, B, E\}$ -setting can be “lifted” to settings with more interfaces. Such a lifting is described by an injective function $\tau : \{A, B, E\} \rightarrow \mathcal{L}$, where we generally assume $\tau(E) = E$. Resources $\mathbf{R} \in \Phi_{\{A, B, E\}}$ are embedded into $\Phi_{\mathcal{L}}$ by providing the A and B -interfaces as $\tau(A)$ and $\tau(B)$ -interfaces and inactive interfaces for all $I \in \mathcal{L} \setminus \tau(\{A, B, E\})$. We denote this resource by $\llbracket \mathbf{R} \rrbracket^{(\tau(A), \tau(B), \tau(E))}$ (we usually only write $\llbracket \mathbf{R} \rrbracket^{(\tau(A), \tau(B))}$). A protocol $\pi = (\pi_1, \pi_2)$ consisting of a pair of converters π_1 for A and π_2 for B becomes $\pi_{\mathcal{L}} = (\pi_I)_{I \in \mathcal{L}}$ with $\pi_{\tau(A)} = \pi_1$, $\pi_{\tau(B)} = \pi_2$, and $\pi_I = \text{id}$ for all $I \notin \tau(\{A, B, E\})$. Security statements transfer from the $\{A, B, E\}$ -setting to the \mathcal{L} -setting since any distinguisher in the \mathcal{L} -setting can be translated into a distinguisher for the $\{A, B, E\}$ -setting by simply emulating the inactive interfaces. (In particular this is a mapping in the above described sense.)

2.5 Discrete Systems

The statements in this paper are statements about *discrete systems*, i.e., systems that communicate by receiving and sending messages. We consider asynchronous systems and do not model time, this means that our definitions currently cover neither timing attacks nor forward secrecy (sometimes called “key corruptions”).

We formalize these systems as random systems [Mau02, Mau13], i.e., families of conditional probability distributions.

Definition 5. Let \mathcal{X} and \mathcal{Y} be sets. An $(\mathcal{X}, \mathcal{Y})$ -random system \mathbf{F} is an infinite sequence of conditional probability distributions $\mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}} : \mathcal{Y} \times \mathcal{X}^i \times \mathcal{Y}^{i-1} \rightarrow [0, 1]$ for $i \geq 1$.

Random systems can be equivalently described by a sequence of conditional probability distributions $\mathbf{p}_{Y_i|X^i}^{\mathbf{F}}$ for $i \geq 1$ that satisfy a compatibility condition, i.e., for any two $i' < i''$ the distributions are consistent for the first i' values.

A game as described in Section 2.3 is an $(\mathcal{X}, \mathcal{Y})$ -system which interacts with its environment by taking inputs X_1, X_2, \dots (considered as moves of the adversary) and answering with outputs Y_1, Y_2, \dots . In addition, after every input it also outputs a bit indicating whether the game has been won. This bit is monotone in the sense that it is initially set to 0 and that, once it has turned to 1 (the game is won), it can not turn back to 0 (even if the game were continued). This motivates the following definition, which captures the notion of game-winning.

Definition 6. For a $(\mathcal{X}, \mathcal{Y} \times \{0, 1\})$ -system \mathbf{F} the binary component A_i of the output (Y_i, A_i) is called a *monotone binary output (MBO)* if $A_i = 1$ implies $A_j = 1$ for $j \geq i$. Such a system \mathbf{F} with MBO is also called a *(discrete) game*.

Any system \mathbf{F} together with a *monotone event sequence (MES)* $\mathcal{A} = (A_1, A_2, \dots)$ defined on \mathbf{F} (within any random experiment, see [Mau02]) can be seen as the game by extending the output of \mathbf{F} by an additional component which turns from 0 to 1 exactly when the MES becomes false. We write $\mathbf{F}^{\mathcal{A}}$ to denote the system with the additional MBO. Conversely, for a $(\mathcal{X}, \mathcal{Y} \times \{0, 1\})$ -system \mathbf{F} (with MBO) we write \mathbf{F}^- to denote the $(\mathcal{X}, \mathcal{Y})$ system where the binary component of the output is discarded.

The adversary (or game winner) and the game are connected via their interfaces; i.e., the adversary specifies the inputs X_i and obtains the outputs Y_i of the game. To formulate “traditional” game-based definitions in this language, the game, often denoted as \mathbf{G} with additional super- and subscripts, allows the adversary \mathbf{A} to issue “oracle queries” via that interface. We usually denote the special monotone output bit as W . For a game \mathbf{G} and an adversary \mathbf{A} , we define the *game-winning probability* after q steps (queries) as

$$\Gamma_q^{\mathbf{A}}(\mathbf{G}) \doteq \mathbb{P}^{\mathbf{A}\mathbf{G}}(W_q = 1).$$

For an adversary \mathbf{A} that halts after (at most) q steps, we write $\Gamma^{\mathbf{A}}(\mathbf{G}) \doteq \Gamma_q^{\mathbf{A}}(\mathbf{G})$.

Two systems with MBOs can be *equivalent as games*, which captures only that they behave equivalently as long as the game is not won.

Definition 7. Two $(\mathcal{X}, \mathcal{Y} \times \{0, 1\})$ -systems with MBO, \mathbf{S} and \mathbf{T} , are *equivalent as games*, denoted $\mathbf{S} \stackrel{g}{\equiv} \mathbf{T}$, if, for $i \geq 1$,

$$\mathbf{p}_{Y^i, A_i=0|X^i}^{\mathbf{S}} = \mathbf{p}_{Y^i, A_i=0|X^i}^{\mathbf{T}}.$$

2.6 Insecure Communication Channels and TLS Fragments

The TLS protocol transmits its data via a TCP connection which transmits a byte stream. At the lowest level, the TLS record protocol then partitions this byte stream into *fragments*, each fragment corresponding to one message sent in the TLS protocol. Technically, each fragment consists of four fields:

1. the *content type* of the transmitted message, which signals whether the message is a `change_cipher_spec` request, an `alert`, a `handshake` message, or `application_data`,
2. the *protocol version*,
3. the *length* of the following payload, and
4. the *payload* message itself.

For simplicity, we model the communication channels that we assume as channels that transmit TLS fragments, ignoring the part of the protocol that converts the fragments into the TCP byte stream and vice versa. The plaintext contained in a TLS plaintext fragment is of length at most $2^{14} = 16384$ bits, and a TLS ciphertext fragment is of length at most $2^{14} + 2048 = 18432$ bits (plus 5 bytes of header). We generally write $\mathcal{B}^{\leq \ell} = \bigcup_{0 \leq i \leq \ell} \{0, 1\}^i$.

3 Constructing the Master Secret

3.1 The Assumed Resources

The resources we assume for the TLS protocol are: an insecure network (obtained by using the TCP/IP protocol in the Internet), a public-key infrastructure which we view as allowing the server to send one message (its public key) authentically to all clients, and a random-oracle resource, outputting consistent random values for user (and adversary) input.

The “Internet” network resource. The insecure network is described as a parallel composition of insecure channels. There are two such channels between each potential client and the server (one in each direction). Technically, the TLS protocol is run via a TCP connection which transmits byte streams; TLS partitions such a byte stream into *TLS fragments* that correspond to the messages sent by higher protocol layers. These fragments are the lowest layer that we consider in our analysis; for us, an insecure communication channel is one that transmits a sequence of TLS fragments.

Insecure channel \rightarrow	
No attacker present:	
Repeatedly:	Upon input a message $m \in \mathcal{B}^{\leq 18432}$ at the <i>A</i> -interface, output m at the <i>B</i> -interface;
Attacker present:	
Repeatedly:	Upon input a message $m \in \mathcal{B}^{\leq 18432}$ at the <i>A</i> -interface, output m at the <i>E</i> -interface.
Repeatedly:	Upon input a message $m \in \mathcal{B}^{\leq 18432}$ at the <i>E</i> -interface, output m at the <i>B</i> -interface.

As TLS builds on TCP, we can envision the insecure channels to be identified by the clients’ TCP sockets, i.e., an address in the set $\mathcal{A}_{\text{TCP}} = \{0, 1\}^{32} \times \{0, 1\}^{16}$. Consequently, the entire network resource, which we denote by **NET**, is defined as the parallel composition of resources $\bigotimes_{C \in \mathcal{A}_{\text{TCP}}} [[- \rightarrow, \leftarrow -]]^{(C, S/C)}$.

The public-key infrastructure. For the case of unilateral authentication (i.e., only the server has a certificate and the clients are not authenticated), we describe a “simple” resource that allows the server to authenticate its public key toward all clients. This task is, in a real protocol, achieved by using certificates obtained from some certification authority.

The resource PKI provides one “client” interface for each $C \in \mathcal{C}$, as well as a “server” interface S , and an “eavesdropper” interface E . The resource is parametrized by a distribution \mathfrak{F} over functions $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ that depends on the scheme used to construct the resource, which might be, as in the case of TLS, X.509. For more details on X.509 see Appendix A.1.

Public-Key Infrastructure Resource $\text{PKI}_{\mathfrak{F}}$

Choose a function $f \leftarrow \mathfrak{F}$.

No attacker present:

- On input a message (**register**, x) at the S -interface with $x \in \{0, 1\}^*$, output $s = f(x)$ at the S -interface.
- On input (**verify**, x, s) at some interface $C \in \mathcal{C}$:
 - If x was input at the S -interface and the response was s , then output 1.
 - Otherwise, output 0.

Attacker present:

- On input a message (**register**, x) at the S -interface with $x \in \{0, 1\}^*$, output $s = f(x)$ at the S -interface. Additionally output (x, s) at the E -interface.
- On input (**verify**, x, s) at some interface $C \in \mathcal{C}$:
 - If x was input at the S -interface and the response was s , then output 1.
 - Otherwise, output 0.

Related papers following the constructive cryptography paradigm, such as [CMT13, MTC13], model the capability of authentically transmitting a single message such as a public key to other parties as an authenticated single-message channel. This simpler formalization is generally preferable because it abstracts from the scheme with which the authentication is achieved; however, in the case of TLS the certificates issued by the certification authority (i.e., the bit strings) are actually used within higher levels of the protocol, in particular, the certificates are included in the computation of the “finished” message.

Uniqueness of certificates. The description of the certification resource $\text{PKI}_{\mathfrak{F}}$ described above formalizes that the encoding of certificates is unique: the verification is successful if and only if *exactly the same* pair (x, s) has occurred previously. This assumption is justified because the X.509 certificates are transmitted in TLS via DER encoding, which guarantees that every data structure has a unique representation.

The random oracle. The random oracle is a resource that can be queried at all interfaces with strings $x \in \{0, 1\}^*$, and for each new string, it responds with a uniform output of length k . If some input was queried (at any interface) before, then the random oracle answers consistently. (If the E -interface is not blocked, the queries at that interface are answered consistently with the queries at the other interfaces.)

Random Oracle RO_k

On input a string $x \in \{0, 1\}^*$ at some interface:

- if x was queried before at some (potentially other) interface, respond with the same value as in that query;
- otherwise, draw a uniformly random $y \in \{0, 1\}^k$ and respond with y .

3.2 Session Naming

In the setting of unilateral key exchange, clients do not *a priori* possess any information that differentiates them from other clients. To prevent protocol sessions from interfering with each other, however, it is necessary to cryptographically bind some unique information to each session. In the TLS protocol, the client’s nonce can serve this purpose. In fact, this nonce can be seen as a “unique name”, as it is—with high probability—unique for honest clients.

3.2.1 Choosing Unique Nonces

Unique names are formalized as a resource that provides to each honest party a unique value. This resource does not have a (non-trivial) E -interface since it merely models the fact that each client can *locally* choose a unique name; there is no requirement that the chosen names need to be known to other parties. For a set \mathcal{N} of names and a set $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ such that $|\mathcal{N}| \geq |\mathcal{C}|$, the resource **NAME** parametrized by an injective function $\rho : \mathcal{C} \rightarrow \rho(\mathcal{C}) \subseteq \mathcal{N}$ assigns to each client $C \in \mathcal{C}$ a unique name $\eta = \rho(C) \in \mathcal{N}$.

Unique name resource NAME_ρ for $\rho : \mathcal{C} \rightarrow \mathcal{N}$

At each interface $C \in \mathcal{C}$, output $\rho(C)$.

Choosing a nonce at random from some distribution \mathbf{N} over the set \mathcal{N} also implements this resource; thus, the resource is constructed without any setup assumptions. In more detail, let rnd be the converter that chooses a nonce $\eta \leftarrow \$\mathbf{N}$ at random and outputs η at the outside interface. The client nonces in TLS consist of a 28 byte random string to which the date and time are prepended, so $\mathcal{N} = \{0, 1\}^{256}$. As the nonce contains 224 bits of randomness, any pair of two clients chooses the same nonce with probability at most $\binom{|\mathcal{C}|}{2} 2^{-224}$, where $|\mathcal{C}|$ is the total number of client sessions (see [MTC13, Lemma 7]). In the following, we use the symbol \mathbf{N} to denote the distribution used in TLS.

3.2.2 Separating Network Sessions

The client’s nonce η_C is sent to the server which uses it to identify that particular client’s session. While the nonces used by honest clients as obtained from the **NAME** resource are unique, the server does not prevent the adversary from starting many sessions with the same nonce. Hence, we index the sessions by pairs $\text{sid} = (\eta_C, e) \in \mathcal{N} \times \mathbb{N}$. Furthermore, the server also chooses a nonce $\eta_{\text{sid}} \leftarrow \\mathbf{N} for this session and sends it to the client. This nonce exchange protocol constructs, from the resources NAME_ρ for some injective mapping $\rho : \mathcal{C} \rightarrow \mathcal{N}$ and **NET**, a resource denoted by **SNET**.

The resource **SNET**, which is described in Figure 3, has interfaces labeled $C \in \mathcal{C}$ for the clients, a server’s interface S which has one sub-interface for each pair (η, e) , where $\eta \in \mathcal{N}$ is the client’s nonce and $e \in [n]$ is a counter indicating how many sessions have been initiated with nonce η , and an adversary’s interface called E . To simplify further construction steps, we rule out collisions for *server* nonces in the **SNET** resource below, in sessions that are associated to

the same client nonce (i.e., $sid = (\eta_C, e)$ and $sid' = (\eta_C, e')$). The reason is that by making the pairs of clients and server nonces unique, we guarantee that any two sessions obtain their keys by querying different parts of the random oracle. Technically, our **SNET** resource will check, when generating a server nonce for a particular client nonce, whether the same server nonce has already been generated for that client nonce. Note that the server nonce has the same structure as the client nonce, thus the security loss is analogous.

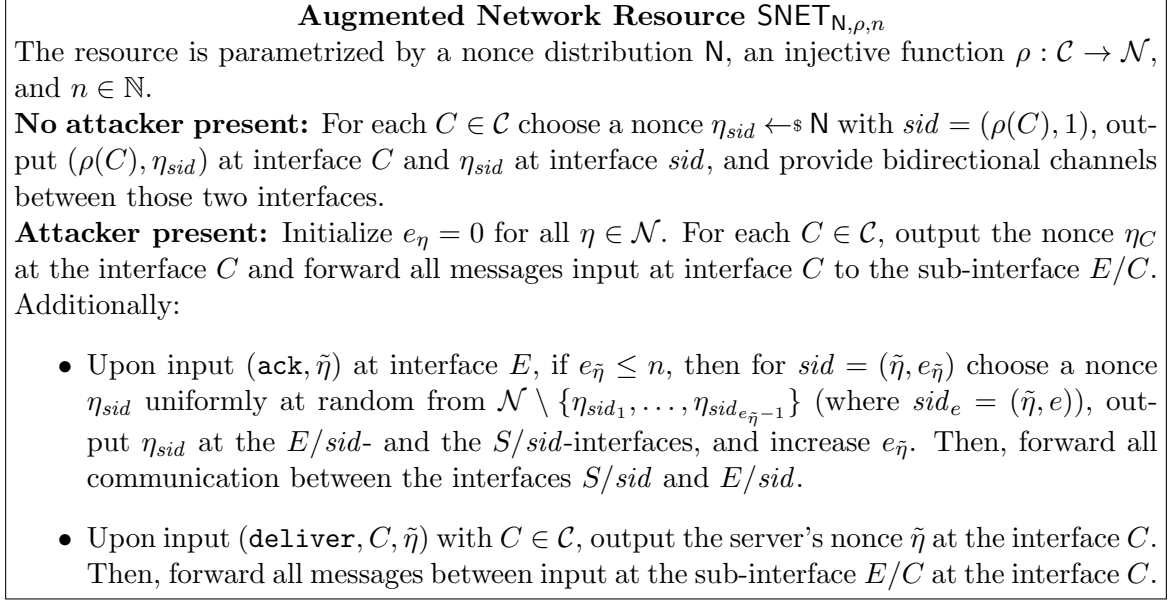


Figure 3: The network resource that additionally outputs nonces.

The resource $\text{SNET}_{\mathbf{N}, \rho, n}$ is constructed (from NAME_ρ and **NET**) by the protocol **(hec, hes_n)** described as follows, where the server's converter **hes_n** is parametrized by the maximum number of sessions it accepts per client nonce.

Client converter hec: (Each such converter connects to an interface $C \in \mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ which has three sub-interfaces: one each for sending messages to, and respectively receiving messages from the server, and a third sub-interface for obtaining the nonce.) Get a nonce $\eta_C \in \mathcal{N}$ from the resource NAME_ρ , output it at the outside interface and send it via the sending sub-interface. Forward all messages input at the outside interface to the first inside sub-interface. Upon receiving $\tilde{\eta}$ at the receiving sub-interface, output $\tilde{\eta}$ at the outside interface. Afterward, forward all messages obtained at the second inside sub-interface to the outside interface.

Server converter hes_n: Initially set $e_\eta = 1$ for each $\eta \in \mathcal{N}$.

- Upon receiving a nonce $\tilde{\eta}_C$ at an inside C -sub-interface for some $C \in \mathcal{A}_{\text{TCP}}$, if $e_{\tilde{\eta}_C} \leq n$, then choose a nonce $\eta_{sid} \leftarrow \$ \mathbf{N}$, send η_{sid} via the C -sub-interface at the inside. Output η_{sid} at outside \tilde{sid} -sub-interface with $\tilde{sid} = (\tilde{\eta}_C, e_{\tilde{\eta}_C})$, and increment $e_{\tilde{\eta}_C}$ by 1.
- Afterward, forward all messages between the outside sub-interfaces associated to \tilde{sid} and the inside C -sub-interfaces.

We prove the following (security) statement.

Lemma 8. Let $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ and let $\rho : \mathcal{C} \rightarrow \mathcal{N}$ be an injective mapping. The protocol $(\text{hec}, \text{hes}_n)$ constructs the resource $\text{SNET}_{\mathbf{N}, \rho, n}$ from the resources NET and NAME_ρ . In more detail, for the simulator σ in the proof:

$$[\text{NET}, \text{NAME}_\rho] \xrightarrow{(\text{hec}, \text{hes}_n), \sigma, (0, \varepsilon)} \text{SNET}_{\mathbf{N}, \rho, n},$$

with $\varepsilon(\mathbf{D}) \doteq \binom{n}{2} \cdot 2^{-224}$ for all distinguishers \mathbf{D} .

Proof.

Availability. Apart from outputting nonces at the server's side, there is one main difference between the real resource, $[\text{NET}, \text{NAME}_\rho]$, and the ideal resource, $\text{SNET}_{\mathbf{N}, \rho, n}$: the server sessions in $\text{SNET}_{\mathbf{N}, \rho, n}$ are denoted by tuples (η, e) , where η is the client nonce and $e \in [n]$ is the index of the session between the server and that particular client; by contrast, the server sessions in NET are denoted as TCP addresses. Since η is unique (by the properties of NET), this difference is strictly nominal, and the server's converter hes_n takes care of it.

Security. We prove the security statement in two steps. In the first step we modify the ideal resource $\text{SNET}_{\mathbf{N}, \rho, n}$ and define a resource \mathbf{H} which behaves exactly like $\text{SNET}_{\mathbf{N}, \rho, n}$ except that it does not check for collisions in server nonce values. We first argue that $\text{SNET}_{\mathbf{N}, \rho, n}$ and \mathbf{H} behave identically, up to a difference of $\binom{n}{2} \cdot 2^{-224}$. Indeed, the two resources behave differently only when, for some nonce $\tilde{\eta}$, the resource \mathbf{H} outputs a duplicate nonce $\eta_{\text{sid}} = \eta_{\text{sid}'}$ for $\text{sid} = (\tilde{\eta}, e) \neq (\tilde{\eta}, e') = \text{sid}'$. A collision in two of at most n sessions occurs with probability $\binom{n}{2} \cdot 2^{-224}$.

Let $\mathbf{R} \doteq \prod_{C \in \mathcal{C}} \text{hec}^C \text{hes}_n^S \perp^E [\text{NET}, \text{NAME}_\rho]$. In the second step, we show \mathbf{R} and $\mathbf{S}' \doteq \sigma^E \mathbf{H}$ are equivalent, for the following σ described below. (Note that intuitively \mathbf{S}' (and indeed the ideal resource $\text{SNET}_{\mathbf{N}, \rho, n}$) has no security properties with respect to the message transmission, e.g. integrity or confidentiality. Indeed, $\text{SNET}_{\mathbf{N}, \rho, n}$ describes a network that exchanges nonces correctly between clients and servers, with the additional property that no two sessions share the same nonces.) The simulator σ behaves as follows:

- *Initialization.* Set $e_\eta = 1$ for all $\eta \in \mathcal{N}$. For each $C \in \mathcal{C}$, simulate $\rho(C)$ as being transmitted on $\llbracket - \rightarrow \rrbracket^{(C, S/C)}$.
- *Output client messages.* All messages obtained via the inside C -sub-interface are output as transmitted via $\llbracket - \rightarrow \rrbracket^{(C, S/C)}$.
- *Delivery of client's nonce.* On input a nonce $\tilde{\eta}$ at the outside sub-interface corresponding to a channel $\llbracket - \rightarrow \rrbracket^{(C, S/C)}$, if $e_{\tilde{\eta}} \leq n$, then input $(\text{ack}, \tilde{\eta})$ at the inside interface, set $\text{sid} = (\tilde{\eta}, e_{\tilde{\eta}})$, and increase $e_{\tilde{\eta}}$. Obtain a nonce at the inside interface, call it η_{sid} , and output η_{sid} at the outside interface as transmitted via $\llbracket \leftarrow - \rrbracket^{(C, S/C)}$. All messages obtained via the inside sid -sub-interface are output as transmitted via $\llbracket \leftarrow - \rrbracket^{(C, S/C)}$, all messages input at the outside sub-interface corresponding to channel $\llbracket - \rightarrow \rrbracket^{(C, S/C)}$ are input at the inside sid -sub-interface.
- *Delivery of the server's nonce.* On input the first message (a nonce $\tilde{\eta}$) at the outside sub-interface corresponding to a channel $\llbracket \leftarrow - \rrbracket^{(C, S/C)}$, input $(\text{deliver}, C, \tilde{\eta})$ at the inside interface. All messages input afterward at the outside sub-interface corresponding to channel $\llbracket \leftarrow - \rrbracket^{(C, S/C)}$ are input at the inside C -sub-interface.

We argue that the simulation is perfect. Firstly, the simulator correctly instantiates the sessions between clients and server. Indeed, in both cases, consistent server sub-interfaces are used because the counters e_{η_C} are updated whenever a nonce η_C is delivered to the server.

The equivalence then follows from three observations:

- For each $C \in \mathcal{C}$, first the nonce η_C is output and then all messages input at the C -interface are output at the E -sub-interface corresponding to $\llbracket - \rightarrow \rrbracket^{(C, S/C)}$ in both cases.
- On input the first message, a nonce $\tilde{\eta}$, at the E -interface of some channel $\llbracket - \rightarrow \rrbracket^{(C, S/C)}$, both the real and the ideal systems will respond with a nonce $\eta_{sid} \leftarrow \$N$ for $sid = (\tilde{\eta}, e_{\tilde{\eta}})$, and also in both cases the output at the interface S/sid is (η_C, η_{sid}) . Afterward, the communication between the interfaces S/sid and the E -interfaces of the channels corresponding to C' is forwarded in both cases.
- On input the first message, a nonce $\tilde{\eta}$, at the E -interface of some channel $\llbracket \leftarrow - \rrbracket^{(C, S/C)}$, both the real and the ideal systems output $\tilde{\eta}$ at the C -interface. Afterward, the communication between the interface C and the E -interfaces of the channels corresponding to C is forwarded in both cases.

This concludes the proof. \square

3.3 Constructing the Shared Key

The next step, following the structure of our outline, is constructing the master secret key. In more detail, we construct (from SNET, PKI, and RO) a resource denoted as $\text{MSK}_{N, \rho, AUX, n}$. Intuitively, this resource represents the result of the core of the handshake itself, as the parties output the master key for the session, from which other keys are derived.

We describe a (slightly) more general resource $\text{KEY}_{N, \rho, AUX, n, \mathcal{K}}$ below, because we will use it in the analysis of TLS 1.3 also for modeling the premaster key (this is not possible in TLS 1.2). This, described in Figure 4, has interfaces $C \in \mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ for the clients, an interface S for the server with sub-interfaces labeled $sid = (\eta, e) \in \mathcal{N} \times [n]$, and an “adversarial” interface E , and is parametrized by the maximum number n of sessions per client nonce. For the set $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$, consider an injective function $\rho : \mathcal{C} \rightarrow \mathcal{N}$, which also parametrizes the resource. The resource is further parametrized by a distribution AUX for the auxiliary information—this is important for specifying the behavior of the resource in the availability case. More generally, this can be a family of distributions to formalize that the auxiliary information in different sessions has some correlated random information. The resource is additionally parametrized by a key space \mathcal{K} .

The master secret key is instantiated with $\mathcal{K} = \{0, 1\}^{384}$, hence with this instantiation $\text{MSK}_{N, \rho, AUX, n} \doteq \text{KEY}_{N, \rho, AUX, n, \mathcal{K}}$. In TLS 1.2, the resource $\text{MSK}_{N, \rho, AUX, n}$ is constructed in one of three ways: by relying on RSA encryption, by means of a Diffie-Hellman key-exchange protocol with a static server key, and by means of an Ephemeral Diffie-Hellman key exchange, with an ephemeral server key. Next, we assess the security in constructing the master secret key resource using TLS-DH, TLS-DHE, and TLS-RSA.

3.3.1 Diffie-Hellman with a Static Server Key

The DH-based protocol consists of two types of converters, one converter dhc attached to each client interface and one converter $\text{dhs}_{\mathcal{G}}$ attached to the server interface. The protocol is run in a DH-group \mathbb{G} of prime order p , generated by some element g which is chosen when the server certificate is generated, according to some distribution which we denote \mathcal{G} . The client converter, dhc , behaves as follows:

1. Obtain the nonces η_C and η_{sid} from $\text{SNET}_{N, \rho, n}$.




Figure 4: The shared key resource.

2. Obtain value $cert$ via $\text{SNET}_{\mathcal{N}, \rho, n}$, parse $cert$ as $((\mathbb{G}, pk_S), s)$.⁴ Here we slightly abuse syntax and write \mathbb{G} to mean a description of the group \mathbb{G} , including the generator g .⁵
3. Verify the server's certificate by querying $(\text{verify}, cert)$ at $\text{PKI}_{\mathbb{G}}$ (if that fails, abort).
4. Choose a secret $u \in \{1, \dots, |\mathbb{G}|\}$ and send $epk_C = g^u$ via $\text{SNET}_{\mathcal{N}, \rho, n}$.
5. Query $pk_S^u | \text{master secret} | \eta_C | \eta_{sid}$ at RO_{384} , receive κ .
6. Output $(\kappa, \eta_C, \eta_{sid}, (epk_C, cert))$.

The server converter, dhs_G , behaves as follows:

0. Sample a group $\mathbb{G} \leftarrow \mathcal{G}$ and generate a DH key pair $(pk_S, sk_S) = (g^v, v)$. Register the

⁴In the specifications of TLS 1.2, if the public key is not extractable from certificates, for client and/or server, a special message can be sent, which includes the public key. Here we restrict to the standard case.

⁵In particular, the X.509 certificate explicitly encodes the prime order p of the Galois field, the prime sub-order q of the group, and the group generator g , of order q [HFPS99].

public key pk_S at $\text{PKI}_{\mathfrak{F}}$ via $(\text{register}, (\mathbb{G}, pk_S))$, obtaining the value $cert = ((\mathbb{G}, pk_S), s)$. For each of the sessions described by a pair (η_C, e) :

1. Upon receiving the nonce η_{sid} at the inside interface, respond with the certificate $cert$.
2. Upon receiving $epk_C \in \mathbb{G}$, query $epk_C | \text{master secret} | \eta_C | \eta_{sid}$ at RO_{384} , receiving $\tilde{\kappa}$.
3. Output $(\tilde{\kappa}, \eta_{sid}, (epk_C, cert))$ at the outer (η_C, e) -sub-interface.

We prove that the Diffie-Hellman-based protocol with static server key indeed constructs the master secret. The distribution AUX associated with the resource $\text{MSK}_{N,\rho,AUX,n}$ produces two values. The first varies for each session and is an element $\tilde{g} \in \mathbb{G}$ (namely the public key epk_C , where \mathbb{G} is sampled according to \mathcal{G}). The second one is fixed for all sessions and is distributed like the certificate $cert$.

Lemma 9. *If the GapDH assumption holds with respect to the distribution \mathcal{G} , then the protocol $(\text{dhc}, \text{dhs}_{\mathcal{G}})$ constructs from $[\text{RO}_{384}, \text{SNET}_{N,\rho,n}, \text{PKI}_{\mathfrak{F}}]$ the resource $\text{MSK}_{N,\rho,AUX,n}$. In more detail, for the simulator σ and the reduction \mathbf{C} in the proof,*

$$[\text{RO}_{384}, \text{SNET}_{N,\rho,n}, \text{PKI}_{\mathfrak{F}}] \xrightarrow{(\text{dhc}, \text{dhs}_{\mathcal{G}}), \sigma, (0, \varepsilon)} \text{MSK}_{N,\rho,AUX,n}$$

such that for all \mathbf{D} , $\varepsilon(\mathbf{D}) = \Gamma^{\mathbf{DC}}(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}})$.

It might appear surprising that we obtain a tight reduction in the construction, given that the security statement applies to a setting with multiple parallel sessions. The reason is that we can exploit the random self-reducibility of the GapDH problem to “inject” the $\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}$ -challenge into every session without changing the overall distribution.

Proof.

Availability. For availability, the ideal $\text{MSK}_{N,\rho,AUX,n}$ resource outputs random values and provides channels between the clients’ and the server’s interfaces that deliver all messages faithfully.

In the real resource the nonces are exchanged in $\text{SNET}_{N,\rho,n}$ —with no attacker present. The certificate is correctly generated for the server’s public key pk_S , and the client’s ephemeral public key epk_C is transmitted correctly. For the output, first note that the key κ , which is output by RO_{384} in the real resource and is picked uniformly at random in the ideal resource, have identical distributions—this is guaranteed by the RO_{384} resource. At the server’s sub-interfaces, the consistent keys is delivered.

Finally, the auxiliary output aux is defined as $(epk_C, cert)$ and has the correct distribution in both cases. Furthermore, once the initial exchanges are done, both resources behave like insecure bidirectional channels in both directions. This verifies the availability condition.

Security. The simulator σ does the following:

- *Initialization.* Initially, set $e_{\eta} = 1$ for all $\eta \in \mathcal{N}$ and define an (initially empty) map $R : \{0, 1\}^* \rightarrow \{0, 1\}^{384}$. Choose a function $f \leftarrow \mathfrak{F}$ and a group $\mathbb{G} \leftarrow \mathcal{G}$, generate the server’s DH key pair $(sk = v, pk = g^v)$ for group \mathbb{G} , compute $s = f(\mathbb{G}, pk)$ and output $cert = ((\mathbb{G}, pk), s)$ at the outside interface to simulate the effect of an input $(\text{register}, pk)$ to $\text{PKI}_{\mathfrak{F}}$.
- *Delivery of the client nonce.* Upon input of the type (ack, η_C) at the outside interface, issue (ack, η_C) at the inside interface⁶ and receive the generated η_{sid} for $sid = (\eta_C, e_{\eta_C})$.

⁶For the analysis, note that in the ideal resource the keys output at the client, server, and E interfaces are all independent of the input nonces; thus, it is not important whether the adversary has forwarded an honestly generated η_C or injected a different $\tilde{\eta}_C$.

Output η_{sid} at the outside interface as coming from SNET, together with the certificate $cert$, and increase e_{η_C} .

- *Delivery of the server nonce.* Upon input $(\text{key-s}, C, \tilde{\eta})$ at the outside interface and after $cert$ is delivered to C via SNET,⁷ generate u , $epk_C = g^u$, and input $(\text{key-c}, C, aux, \tilde{\eta})$ at the inside interface (where $aux = (cert, epk_C)$). Simulate sending the DH element epk_C from C via $\text{SNET}_{N,\rho,n}$, and register $(\tilde{\eta}, epk_C)$.
- *Delivery of DH element.* If some group element $\tilde{epk}_C \in \mathbb{G}$ is input at the outside sub-interface $sid = (\eta, e)$ with $e \geq e_\eta$, then:
 - If \tilde{epk}_C corresponds to a recorded pair $(\eta_{sid}, \tilde{epk}_C)$ then issue $(\text{deliver}, \eta, e, aux)$ at the inside interface, with $aux = (cert, epk_C)$;
 - Otherwise, input $(\text{inject}, \eta, e, aux, \kappa)$ at the inside interface, with $aux = (cert, \tilde{epk}_C)$, and with the value κ computed as follows: if $\tilde{epk}_C^v | \text{master secret} | \eta | \eta_{sid}$ is defined in R , then use $R(\tilde{epk}_C^v | \text{master secret} | \eta | \eta_{sid})$; otherwise define it as a freshly sampled value $R(\tilde{epk}_C^v | \text{master secret} | \eta | \eta_{sid}) \leftarrow_{\$} \{0, 1\}^{384}$ and use that.
- *Simulation of random oracle.* Whenever the adversary E or either of the honest parties are supposed to query the RO_{384} resource, return a random element $R(x) \leftarrow_{\$} \{0, 1\}^{384}$, unless the same input was queried before (in that case, return the value that was previously returned).

Denote $\mathbf{R} := \prod_{C \in \mathcal{C}} \text{dhc}^C \text{dhs}_G^S[\text{RO}_{384}, \text{SNET}_{N,\rho,n}, \text{PKI}_{\mathfrak{S}}]$ and let $\mathbf{S} := \sigma^E \text{MSK}_{N,\rho,AUX,n}$. We first analyze the behavior of \mathbf{R} and \mathbf{S} in the initial queries.

- The exchange of the server’s nonces in the real and the ideal resource (the former using SNET for it) is identical, the output at the E -interface having exactly the same distributions.
- The value $cert$ has the same distribution in both cases.
- Injecting a value \tilde{epk}_C in a session between client C and the interface $sid = (\eta, e)$, such that the client C has already sent $epk_C \neq \tilde{epk}_C$ and computed its output, does not cause discrepancies. In particular, as the keys are injected by the simulator and the random oracle is simulated consistently, the simulation in this case is perfect, irrespective of any queries the adversary makes to the random oracle.

Intuitively, the discrepancy between the two resources appears when the distinguisher expects a different output from the random oracle resource RO_{384} (in the real world) than the output generated by the simulator (in the ideal world). This occurs (only) in sessions where the key κ output by the client or server is obtained from RO_{384} via a query that is also asked at RO_{384} by the adversary, whereas in the ideal world, the output at the client’s or server’s interface is picked uniformly at random by $\text{MSK}_{N,\rho,AUX,n}$ and the simulator cannot simulate queries to RO_{384} consistently.

In the following, we show a reduction such that if such an event happens for *some* session, this reduction outputs a solution to the $\mathbf{G}_G^{\text{GapDH}}$ game (see Appendix C.2)—in particular, the distinguisher must be able to find the input value \tilde{epk}_C^v by only knowing $epk_C = g^u$ and g^v . To simulate the environment for the distinguisher, the reduction will make “DDH-queries” at $\mathbf{G}_G^{\text{GapDH}}$, which enables it to simulate the random oracle consistently.

⁷Note that we simplify the certificate verification step in our analysis, confining ourselves to simply outputting a bit, indicating validity/non-validity. In TLS, this is done in a more thorough manner, and several alert messages are generated, depending on where the verification failed.

We describe a reduction **C** that obtains $(\mathbb{G}, g, g_1, g_2)$ at the inside interface and starts simulating the setup using the group description (\mathbb{G}, g) and g_1 as the server's public key (thus for the simulated instance it holds that $g^v = g_1$), and otherwise behaves similar to **S**. To simulate a group element sent by the client C_i , **C** chooses a value $r_i \in [|\mathbb{G}|]$ uniformly at random and simulates sending $\tilde{g}_i = g_2 \cdot g^{r_i}$.

As the reduction does not know the server's secret, i.e. v , it has to make DDH-queries at $\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}$ to answer random oracle queries consistently. In more detail, the simulation of queries goes as follows. For administrative purposes, the reduction keeps track of what the distinguisher knows about the random oracle by updating two lists, initially empty. One list, which we denote as \mathcal{H}_H records queries of the form (x, η, η_{sid}, y) (the value y is the output to honest parties to input of the form $x|\text{master secret}|\eta|\eta_{sid}$). The second list is denoted \mathcal{H}_M and records malicious deliveries of \tilde{epk}_C values by the distinguisher to the simulated server; the entries are of the form $(\tilde{epk}_C, \eta, \eta_{sid}, y)$, where y is the output returned by the reductions when simulating RO_{384} for the maliciously injected input \tilde{epk}_C . The queries are answered by the reduction as follows:

- Whenever the random oracle is queried on input (x, η, η_{sid}) , first check if there exists an entry $(x, \eta, \eta_{sid}, y) \in \mathcal{H}_H$ and if so, output y ; else check if $\text{DDH}(g_1, g_2, x \cdot g_1^{-r_i}) = 1$ for some $i \in \{1, \dots, |\mathcal{C}|\}$ and if so, output $x \cdot g_1^{-r_i}$ as solution to $\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}$. Else check for all entries $(\tilde{epk}_C, \eta, \eta_{sid}, y) \in \mathcal{H}_M$ and return y if $\text{DDH}(\tilde{epk}_C, g_1, x) = 1$. Otherwise, respond with a random y and update the list $\mathcal{H}_H \leftarrow \mathcal{H}_H \cup \{(x, \eta, \eta_{sid}, y)\}$.
- When the adversary delivers some group element \tilde{epk}_C to the server session sid , the reduction needs to simulate a random oracle query on input $(\tilde{epk}_C^v, \eta, \eta_{sid})$, where $v = \log_g g_1$. As the reduction does not know v , it first checks if there is an entry $(\tilde{epk}_C, \eta, \eta_{sid}, y) \in \mathcal{H}_M$; if such an entry exists, the reduction returns y ; else, it checks if there exists an input of the form $(x, \eta, \eta_{sid}, y) \in \mathcal{H}_H$ such that $\text{DDH}(\tilde{epk}_C, g_1, x) = 1$ and if such a value exists, it returns it and updates $\mathcal{H}_M \leftarrow \mathcal{H}_M \cup \{(\tilde{epk}_C, \eta, \eta_{sid}, y)\}$. If no such value exists then the reduction returns a random value y , and also updates $\mathcal{H}_M \leftarrow \mathcal{H}_M \cup \{(\tilde{epk}_C, \eta, \eta_{sid}, y)\}$.

We define MBOs $\mathcal{E}^i = (E_1^i, E_2^i, \dots)$ on the systems **R**, **S**, and **C** $(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}})^-$ as follows: E_q becomes 1 if after q queries there has been a query $(pmk|\text{master secret}|\eta_C|\eta_{sid})$ at the random oracle, such that $C = C_i$, $sid = (\eta_C, j)$ for some $j \in [n]$, and pmk is the DH element corresponding g_1 and \tilde{g}_i . We then set $\mathcal{E} = \bigvee_{i=1}^n \mathcal{E}^i$. Note that for $g_1 = g^u$, $g_2 = g^v$ (so $\tilde{g}_i = g^{v+r_i}$) and $pmk = g^{u(v+r_1)} = g^{uv} \cdot g^{ur_1} = g^{uv} \cdot g_1^{r_1}$, the reduction **C** always obtains the correct $g^{uv} = pmk \cdot g_1^{-r_1}$.

Now it holds that $\mathbf{R}^{\mathcal{E}} \stackrel{g}{=} \mathbf{S}^{\mathcal{E}} \stackrel{g}{=} \mathbf{C}^{\mathcal{E}} (\mathbf{G}_{\mathcal{G}}^{\text{GapDH}})^-$, in particular all clients' group elements are uniformly distributed by their definition in **dhc**, σ , and **C**. Using Lemma 26 together with the fact that **C** wins the game $\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}$ whenever the output \mathcal{E} is provoked, formally $\Gamma^{\text{DC}}(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}) = \Gamma^{\text{D}}(\mathbf{C}\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}) \geq \Gamma^{\text{D}}(\mathbf{C}^{\mathcal{E}} (\mathbf{G}_{\mathcal{G}}^{\text{GapDH}})^-)$, concludes the proof. \square

3.3.2 Diffie-Hellman with an Ephemeral Server Key

The construction of the master secret key resource using the ephemeral Diffie-Hellman mode can be proven in two constructive steps. First, we use the signature scheme to construct a network $\succsim \bullet$ from **SNET** and **PKI**; this network allows the server to send one message in each session in an authenticated way. Second, we complete the analysis of the TLS-DHE mode by using the resource $\succsim \bullet$ to transmit the group parameters and the server's ephemeral public key.

The authenticated transmission resource. The resource $\succsim \bullet$ has (client) interfaces $C \in \mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$, a (server) interface S with sub-interfaces $(\eta, e) \in \mathcal{N} \times \mathbb{N}$ with $e \in \mathbb{N}$, and an (eavesdropper) interface E with sub-interfaces $\mathcal{A}_{\text{TCP}} \cup (\mathcal{N} \times \mathbb{N})$, and is parametrized by an injection $\rho : \mathcal{C} \rightarrow \mathcal{N}$, a distribution \mathfrak{F} over functions $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, and a signature scheme $SIG = (\text{gen}, \text{sign}, \text{vrf})$ as defined in Appendix B.1. The resource is described in detail in Figure 5.

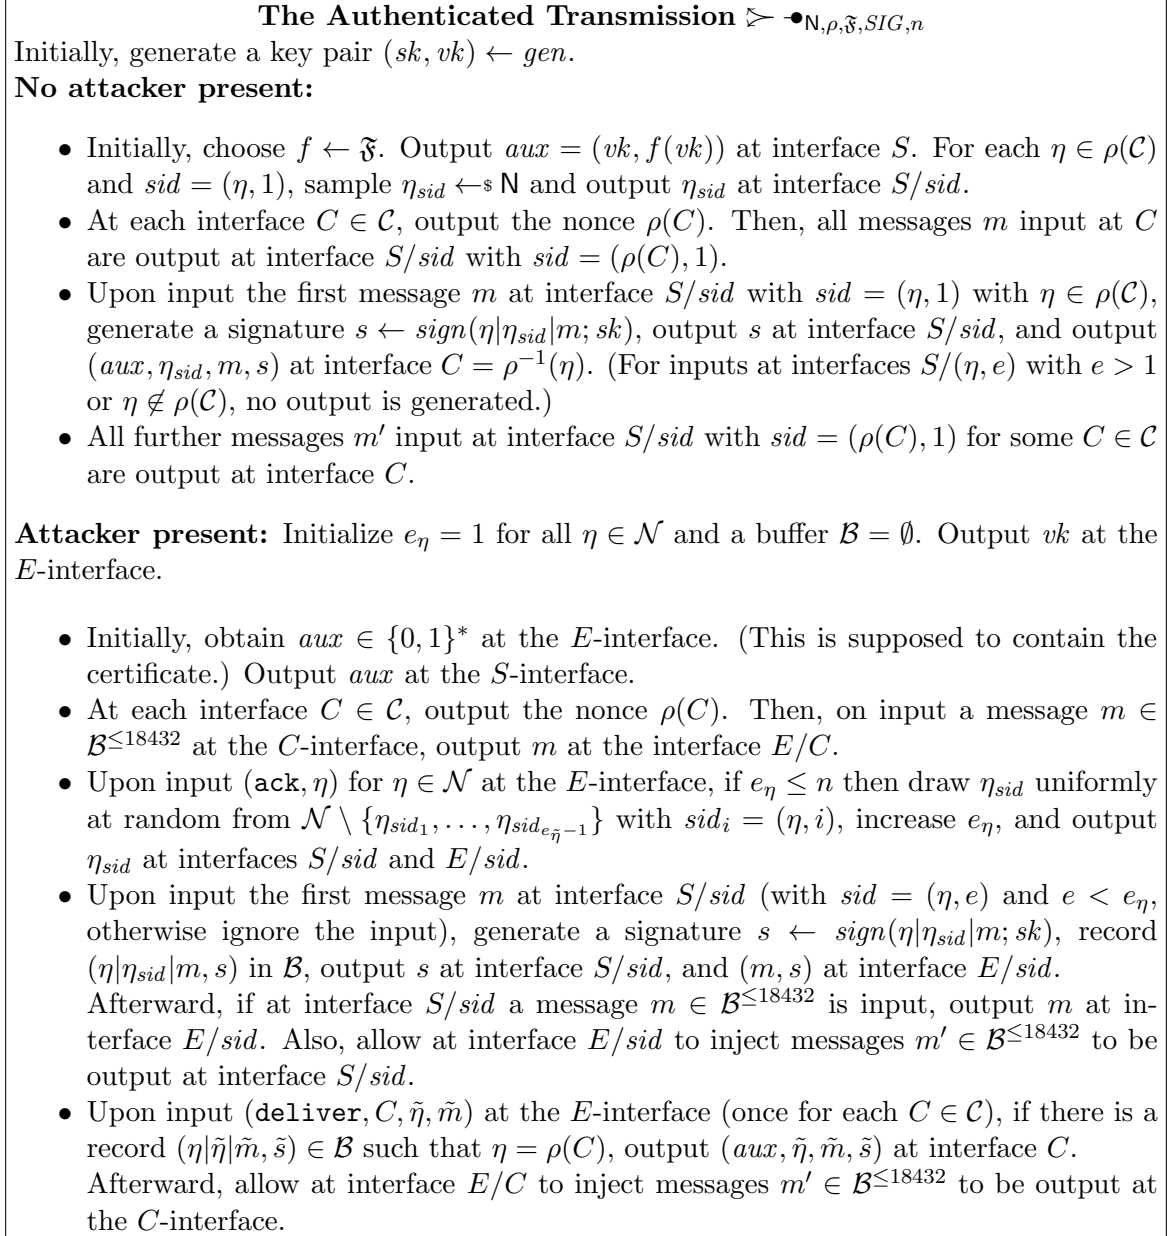


Figure 5: The network allowing the server to transmit, to each client, one message authentically. TLS 1.2 variant.

The following protocol constructs $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n}$ from $\text{SNET}_{\mathbf{N}, \rho, n}$ and $\text{PKI}_{\mathfrak{F}}$: The client's converter vrf is based on the signature scheme SIG and behaves as follows:

0. Obtaining the nonce $\eta_C \in \mathcal{N}$ from $\text{SNET}_{\mathbf{N}, \rho, n}$, output η_C at the outside interface. Forward

all messages from the outside to the inside interface. (Messages sent from the client to the server.)

1. Obtain the server's nonce $\eta_{sid} \in \mathcal{N}$ from $\text{SNET}_{\mathbf{N}, \rho, n}$.
2. Obtain the first message $cert$ from $\text{SNET}_{\mathbf{N}, \rho, n}$. Query $(\text{verify}, cert)$ at $\text{PKI}_{\mathfrak{F}}$; abort if the verification fails or if $cert$ is not a well-formed certificate $cert = (vk, f(vk))$.
3. Obtain the second message m' from $\text{SNET}_{\mathbf{N}, \rho, n}$ and parse m' as (m, s) (abort if that is impossible). If $\text{vrf}(\eta_C | \eta_{sid} | m, s; vk) = 1$, then output $(cert, \eta_C, \eta_{sid}, m, s)$ at the outside interface. (Otherwise abort.)
4. Forward all further messages from the inside to the outside interface. (Messages sent from the server to the client.)

The server's converter sgn provides at the outside "sessions" for all $sid = (\eta_C, e) \in \mathcal{N} \times [n]$ and behaves as follows:

0. Compute $(sk, vk) \leftarrow \text{gen}$. Input vk at $\text{PKI}_{\mathfrak{F}}$, obtaining a response s , and set $cert = (vk, s)$. Output $cert$ at the outside interface (as auxiliary information).
- For each session $sid = (\eta_C, e)$ —i.e., the inputs/outputs at $\text{SNET}_{\mathbf{N}, \rho, n}$ are at the corresponding inside sub-interface sid , and the inputs/outputs at the outside interface are at sub-interface sid —do:
1. Receiving a nonce η_{sid} from $\text{SNET}_{\mathbf{N}, \rho, n}$, output η_{sid} at the outside and send $cert$ via $\text{SNET}_{\mathbf{N}, \rho, n}$.
 2. Obtaining a message m at the outside, compute $s \leftarrow \text{sign}(\eta_C | \eta_{sid} | m; sk)$ and send (m, s) via $\text{SNET}_{\mathbf{N}, \rho, n}$. Output s at the outside.
 3. Forward messages between inside and outside.

Lemma 10. *The protocol (vrf, sgn) for a particular signature scheme $SIG = (\text{gen}, \text{sign}, \text{vrf})$ constructs $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n}$ from $\text{SNET}_{\mathbf{N}, \rho, n}$ and $\text{PKI}_{\mathfrak{F}}$, if the signature scheme SIG is unforgeable. In more detail, for the simulator σ and the reduction \mathbf{C} described in the proof,*

$$[\text{SNET}_{\mathbf{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}] \xrightarrow{(\text{vrf}, \text{sgn}), \sigma, (0, \varepsilon)} \succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n},$$

with $\varepsilon(\mathbf{D}) \doteq \Gamma^{\mathbf{DC}}(\mathbf{G}^{\text{uf-cma}})$ for all distinguishers \mathbf{D} .

Proof.

Availability. We first verify the availability condition. The keys (sk, vk) and the function f have the same distribution in the real and the ideal case,⁸ the same holds for the auxiliary information $aux = (vk, f(vk))$ output at interface S . Each interface S/sid with $sid = (\eta_C, 1)$ and $\rho(C) = \eta_C$ outputs a random nonce η_{sid} also in both cases. On input the first message m at interface S/sid , the same interface outputs a (valid) signature s for m , and the client's interface C outputs $(aux, \eta_C, \eta_{sid}, m, s)$ (since all messages can be parsed correctly, the verification of the server's certificate at $\text{PKI}_{\mathfrak{F}}$ succeeds, and the verification of the signature s succeeds by the correctness of SIG). Afterward, messages input at interface C are output at interface S/sid and vice versa—in the real case these are simply transmitted via $\text{SNET}_{\mathbf{N}, \rho, n}$.

Security. For proving the security condition, we consider the following simulator σ :

⁸In the composed scheme, we let \mathfrak{F} be the distribution that one gets by generating a key pair according to the signature scheme used by the assumed public-key infrastructure, and then for a parameter $x \in \{0, 1\}^*$ the result of f is the pair (x, s) with $s \leftarrow \text{sign}(x; sk)$.

- Initially, σ obtains the public key vk at the inside interface, chooses a function $f \leftarrow \mathfrak{F}$, and computes $cert = (vk, f(vk))$. Also, σ internally initializes $e_\eta = 1$ for all $\eta \in \mathcal{N}$, and inputs $cert$ as auxiliary information at the inside interface, and as output of $\text{PKI}_{\mathfrak{F}}$ at the outside interface.
- Upon input (ack, η) for a nonce $\eta \in \mathcal{N}$ at the outside interface, set $sid = (\eta, e_\eta)$. If $e_\eta \leq n$, then input (ack, η) at the inside interface. Obtain the nonce η_{sid} at the inside interface, output η_{sid} as response of $\text{SNET}_{\mathbf{N}, \rho, n}$ at the outside interface, and increase e_η . Also, output $cert$ at the outside interface as being transmitted via $\text{SNET}_{\mathbf{N}, \rho, n}$ in session sid .
- Upon input $(\text{deliver}, C, \tilde{\eta})$ at the outside sub-interface corresponding to $\text{SNET}_{\mathbf{N}, \rho, n}$, if this is the first such query for C , record the nonce as $\eta_{sid} := \tilde{\eta}$ for $sid = (\rho(\eta), *)$.
- When obtaining an output (m, s) at the inside sub-interface sid with $sid = (\eta, e)$, output (m, s) as being transmitted on $\text{SNET}_{\mathbf{N}, \rho, n}$ as the second message via the corresponding sub-interface. Mark sid as “active”.
- When receiving at the outside sub-interface corresponding to $\text{SNET}_{\mathbf{N}, \rho, n}$ the first message \tilde{m}_1 to a client C , record C as “failed” if $\tilde{m}_1 \neq cert$.
- When receiving at the outside sub-interface corresponding to $\text{SNET}_{\mathbf{N}, \rho, n}$ the second message $\tilde{m}_2 = (m', s')$ to a client C , if C is not marked as “failed” and, with $sid = (\eta_C, *)$, $\text{vrf}(\eta_C | \eta_{sid} | m', s'; vk) = 1$, then input $(\text{deliver}, C, \eta_{sid}, m')$ at the inside interface⁹ and mark C as “active”.
- Communication is forwarded: for clients $C \in \mathcal{C}$ marked as “active,” messages $m \in \mathcal{M}$ obtained at sub-interface C at the inside are output at the outside as being sent by C via $\text{SNET}_{\mathbf{N}, \rho, n}$, and messages $m' \in \mathcal{M}$ obtained at the outside as input to $\text{SNET}_{\mathbf{N}, \rho, n}$ directed to C are input at sub-interface C at the inside. The analogous behavior applies to the (server-bound) “active” sub-interfaces described by $sid = (\eta, e)$.

Moreover, we use the reduction \mathbf{C} that connects with the inside interface to the game $\mathbf{G}^{\text{uf-cma}}$ and provides at the outside interface an emulation of $\mathbf{R} := \prod_{C \in \mathcal{C}} \text{vrf}^C \text{sgn}^S[\text{SNET}_{\mathbf{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}]$, using the key pair (sk, vk) obtained from $\mathbf{G}^{\text{uf-cma}}$. Roughly, \mathbf{C} emulates \mathbf{R} , but computes the certificate $cert = (vk, f(vk))$ using the verification key vk obtained from $\mathbf{G}^{\text{uf-cma}}$, and the signatures $s \leftarrow \text{sign}(\eta_C | \eta_{sid} | m; sk)$ in the converter sgn using signing queries to $\mathbf{G}^{\text{uf-cma}}$. A “forgery” in the emulated execution can then be used to win the game $\mathbf{G}^{\text{uf-cma}}$ (for the exact definition of forgery see the MBO below; the forgery for $\mathbf{G}^{\text{uf-cma}}$ is achieved by concatenating the corresponding nonces and the message).

We define the following MBO \mathcal{E} on the systems \mathbf{R} , $\mathbf{S} := \sigma^{E \curvearrowright} \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$, and $\mathbf{C}\mathbf{G}^{\text{uf-cma}}$ as the following “no forgery”-event: it becomes 1 once there is an input (\tilde{m}, \tilde{s}) at the E/sid -interface with $sid = (\eta_C, e)$ such that $\text{vrf}(\eta_C | \eta_{sid} | \tilde{m}, \tilde{s}; vk) = 1$ —such that η_{sid} was generated as a response in the S -session sid —unless there was an output (\tilde{m}, \tilde{s}') at some interface E/sid' with $sid' = (\eta, e')$ and $\eta_{sid'} = \eta_{sid}$ before. Then, proving “game equivalence” and using Lemma 26 allows to conclude that the condition is fulfilled.

The equivalence can be seen as follows (we use the counter variables e_η in the same way they are defined in the systems, i.e., counting the number of sessions that have been initiated with nonce η):

- Initially, both systems output the certificate $cert = (vk, f(vk))$, with vk and f chosen according to the same distributions, as auxiliary information at the S -interface, and as an output corresponding to $\text{PKI}_{\mathfrak{F}}$ at the E -interface. Additionally, for each $C \in \mathcal{C}$, both

⁹This has an effect only if the corresponding message has been sent by the server before; if the signature is forged, the real and ideal systems behave differently.

systems output the client's nonce η_C at the C -interface and forward messages from the C -interface to the E/C -sub-interface.

- Upon input (ack, η) at the outside E -interface corresponding to $\text{SNET}_{N,\rho,n}$, the system **R** outputs a nonce η_{sid} at the E -interface of $\text{SNET}_{N,\rho,n}$ and the S/sid for $sid = (\eta, e_\eta)$, as well as $cert$ being sent via $\text{SNET}_{N,\rho,n}$ from the server-session sid to the client C (see the description of sgn). In **S**, the nonce is generated according to the same distribution and output at the same interfaces; the message $cert$ is simulated correctly.
- Upon input $(\text{deliver}, C, \tilde{\eta})$ at the outside E -interface corresponding to $\text{SNET}_{N,\rho,n}$, there is no immediate output (neither for **R** nor for **S**).
- Upon delivering the first message \tilde{m}_1 via $\text{SNET}_{N,\rho,n}$ to a client C , if $\tilde{m}_1 \neq cert$ then, in **R**, vrf aborts since either \tilde{m}_1 is not a well-formed certificate or the verification at $\text{PKI}_{\mathfrak{F}}$ fails. In **S**, σ marks C as “failed” in the same cases. There is no output, neither in **R** nor in **S**.
- Upon the first input m at the outside S/sid -interface with $sid = (\eta, e)$, system **R** generates a signature s for $\eta_C|\eta_{sid}|m$ using the key sk (within sgn), outputs s at the same sub-interface and the pair (m, s) at the E -interface of $\text{SNET}_{N,\rho,n}$. Within **S**, the signature is computed analogously by $\succsim \bullet_{N,\rho,\mathfrak{F},SIG,n}$, also output at S/sid , and the message/signature pair is output at the E -interface via σ .
- Upon delivering the second message \tilde{m}_2 via $\text{SNET}_{N,\rho,n}$ to a client C , if \tilde{m}_2 can be parsed as a pair (m, s) such that, with $sid = (\eta_C, *)$, $\text{vrf}(\eta_C|\eta_{sid}|m, s; vk) = 1$, then in **R** $(cert, \eta_{sid}, m, s)$ is output at the C -interface (and nothing if vrf aborted earlier or because the verification failed). By the definition of \mathcal{E} , this means that there was an output (m, s') at some interface E/sid' with $sid' = (\eta_C, e)$, $\eta_C = \rho(C)$, and $\eta_{sid'} = \eta_{sid}$ (where η_{sid} was input via $(\text{deliver}, C, \eta_{sid})$). In **S**, the simulator σ also checks whether the message \tilde{m}_2 can be parsed correctly, the client has not “failed,” and the signature verifies for the message extended by prepending the nonces η_C and η_{sid} with $sid = (\eta_C, *)$, where the latter nonce was delivered to C as the server's nonce before. The MBO assures that there exists a corresponding record in the buffer \mathcal{B} , which means that the output of **R** and **S** is consistent if no forgery occurs, i.e., the systems are equivalent as games.
- After a certain session (either $C \in \mathcal{C}$ or sid at S) has been initialized, messages input there are output at the corresponding sub-interface of the E -interface and vice versa. This is consistent in both **R** and **S**.

As the same arguments (with the exception that the signatures are obtained from $\mathbf{G}^{\text{uf-cma}}$ but have the same distribution) hold for the distribution in the case $\mathbf{CG}^{\text{uf-cma}}$, and each violation of \mathcal{E} can be used to win $\mathbf{G}^{\text{uf-cma}}$, this concludes the proof. \square

Constructing the key. The subsequent construction step is then achieved by the protocol $(\text{dhec}, \text{dhes}_G)$, in which the server chooses for each session a (potentially fresh) Diffie-Hellman group and element, which are sent as an authenticated message via $\succsim \bullet$. We denote the distribution over groups that the server uses by \mathcal{G} ; the only restriction implied by the TLS standard is that the group specified as \mathbf{Z}_p^\times , where p is represented by at most 65535 bits [DR08, page 51].¹⁰ We write dhes_G wherever we want to make the distribution \mathcal{G} used by the server converter explicit.

The distribution AUX in this case consists of two parts. The first part is the same for all sessions and consists of the certificate (i.e., depends on \mathfrak{F} and SIG of $\succsim \bullet_{N,\rho,\mathfrak{F},SIG,n}$). The distribution is described by $(sk, vk) \leftarrow \text{gen}, f \leftarrow_{\$} \mathfrak{F}$, and then $cert = (vk, f(vk))$. The second part is chosen independently for each session by doing the same process as in the protocol:

¹⁰We assume that the server picks a safe prime of appropriate size.

choose a prime $p \in \mathbb{N}$ and a generator $g \in \mathbf{Z}_p^\times$ according to the distribution \mathcal{G} , and choose two group elements $g_1, g_2 \leftarrow \mathbf{Z}_p^\times$ uniformly at random.

The client's converter dhec obtains η_C and (aux, η_{sid}, m, s) at the inside interface, and then:

- Parse the message as $p|g|g' = m$ (abort if impossible).
- Choose $u \leftarrow \{1, \dots, q\}$ (with $q = |\mathbf{Z}_p^\times|$) and input g^u at the inside interface.
- Query $g'^u | \text{master secret} | \eta_C | \eta_{sid}$ at RO_{384} , in order to obtain a key $\kappa \in \{0, 1\}^{384}$. Output $(\kappa, \eta_C, \eta_{sid}, aux | m | s | g^u)$.
- Forward the following communication between the inside and the outside interfaces.

The server's converter dhes_G connects to the S -interface of $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$. Both the inside and outside interfaces have sub-interfaces $sid = (\eta_C, e) \in \mathcal{N} \times [n]$. The converter behaves as follows:

- Initially, receive aux on the inside interface.
- Upon obtaining a nonce η_{sid} at sub-interface $sid = (\eta_C, e)$, choose a modulus $p \in \mathbb{N}$ and a generator $g \in \mathbf{Z}_p^\times$ according to the distribution \mathcal{G} . Also, choose an exponent $v \leftarrow \{1, \dots, |\mathbf{Z}_p^\times|\}$. Send $m = p|g|g^v$ via the inside sub-interface sid . Obtain the signature s at the inside interface in return.
- Upon receiving a group element \tilde{g} in an (active) session $sid = (\eta_C, e)$ at the inside interface, query $\tilde{g}^v | \text{master secret} | \eta_C | \eta_{sid}$ at RO_{384} , call the result κ . Output $(\kappa, \eta_{sid}, aux | m | s | \tilde{g})$ at the outside sub-interface sid .
- Forward the following communication between the corresponding sub-interfaces sid of the inside and the outside interface.

The described protocol indeed constructs the master secret key resource from the network $\succsim \bullet$ and the random oracle RO_{384} under the assumption that the GapDH assumption holds with respect to the distribution \mathcal{G} .

Lemma 11. *The protocol $(\text{dhec}, \text{dhes}_G)$ constructs from the assumed resources $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$ and RO_{384} the resource $\text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$, under the GapDH assumption for \mathcal{G} . More formally, for the simulator σ and the reduction \mathbf{C} described in the proof,*

$$[\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}, \text{RO}_{384}] \xrightarrow{(\text{dhec}, \text{dhes}_G), \sigma, (0, \varepsilon)} \text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n},$$

with $\varepsilon(\mathbf{D}) = n \cdot |\mathcal{C}| \cdot \Gamma^{\text{DC}} \left(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}} \right)$ for all distinguishers \mathbf{D} .

Proof.

Availability. We first prove the availability condition. The “ideal” system $\perp^E \text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ chooses, for each session $sid = (\eta_C, 1)$ with $\rho(C) = \eta_C$, a nonce η_{sid} , a key κ_C and auxiliary information aux_C , and outputs $(\kappa_C, \eta_{sid}, aux_C)$, and afterward forwards communication between C and S/sid .

The distribution in the case $\prod_{C \in \mathcal{C}} \text{dhec}^C \text{dhes}_G^S \perp^E [\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}, \text{RO}_{384}]$ is as follows. The key κ_C for each client C is a uniformly random 384-bit string (the random oracle RO_{384} is queried at distinct places as η_C is distinct by the assumption on ρ), the nonces η_{sid} for $sid = (\eta_C, 1)$ are distributed according to \mathbf{N} , and the auxiliary information aux_C has exactly the same distribution as well. Moreover, after at both interfaces C and S/sid the above information is output, the resource behaves as a bidirectional channel between those interfaces. This verifies the validity of the condition.

Security. The simulator basically needs to take care of the $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$'s E -interface (beginning with choosing some good aux), and then needs to simulate the DH exchange, i.e., the group and elements generated by the server as well as the element generated by the client. The simulator σ behaves as follows:

- Throughout, σ keeps counters e_η for each $\eta \in \mathcal{N}$ in the usual way (i.e., increase e_η whenever the nonce η is delivered to the server). Also, σ keeps a map $R : \{0, 1\}^* \rightarrow \{0, 1\}^{384}$ which is initially empty.
- Initially, obtain a string $\text{cert} \in \{0, 1\}^*$ at the outside interface (this is needed for the auxiliary information).
- Upon input (ack, η) at the outside, input (ack, η) at the inside and obtain as response η_{sid} for $\text{sid} = (\eta, e_\eta)$. Output η_{sid} at the outside interface. Also, choose a modulus $p_{\text{sid}} \in \mathbb{N}$ and a generator $g_{\text{sid}} \in \mathbf{Z}_{p_{\text{sid}}}^\times$ according to \mathcal{G} , as well as a value $v \leftarrow \{1, \dots, q_{\text{sid}}\}$ for $q_{\text{sid}} = |\mathbf{Z}_{p_{\text{sid}}}^\times|$, compute $\tilde{g}_{\text{sid}} = g_{\text{sid}}^v$ and $s \leftarrow \text{sign}(\eta | \eta_{\text{sid}} | p_{\text{sid}} | g_{\text{sid}} | \tilde{g}_{\text{sid}}; sk)$, and output $(p_{\text{sid}} | g_{\text{sid}} | \tilde{g}_{\text{sid}}, s)$ at the outside sub-interface sid .
- Upon input $(\text{deliver}, C, \tilde{\eta}, \tilde{m})$ at the outside, if (ack, η_C) was input before, and there is a session $\text{sid} = (\eta_C, e)$ with $\tilde{\eta} = \eta_{\text{sid}}$ and $\tilde{m} = p_{\text{sid}} | g_{\text{sid}} | \tilde{g}_{\text{sid}}$, then input $(\text{key-c}, C, \text{aux}, \eta_{\text{sid}})$ with $\text{aux} = \text{cert} | p_{\text{sid}} | g_{\text{sid}} | \tilde{g}_{\text{sid}} | \bar{g}_{\text{sid}}$ where $\bar{g}_{\text{sid}} = g^u$ for a uniformly random $u \in \{1, \dots, q_{\text{sid}}\}$. Simulate \bar{g}_{sid} as the first message transmitted from C to sid . (If any check fails, output nothing—note that we simplify the protocol and do not handle error messages.)
- If a group element g'_{sid} is delivered (i.e., input at the outside interface) to some instance sid (and g'_{sid} is a valid group element in the group with modulus p_{sid}):
 - if $g'_{\text{sid}} = \bar{g}_{\text{sid}}$, then input $(\text{key-s}, \eta_C, e, \text{aux})$ and $\text{aux} = \text{cert} | p_{\text{sid}} | g_{\text{sid}} | \tilde{g}_{\text{sid}} | \bar{g}_{\text{sid}}$ (with $\text{sid} = (\eta_C, e)$) at the inside interface;
 - otherwise, for $x = g'_{\text{sid}}^v | \text{master secret} | \eta_C | \eta_{\text{sid}}$, if $R(x)$ is undefined, initialize it as $R(x) \leftarrow \{0, 1\}^{384}$ and input $(\text{inject}, \eta_C, e, \text{aux}, R(x))$ at the inside interface with $\text{sid} = (\eta_C, e)$ and $\text{aux} = \text{cert} | p_{\text{sid}} | g_{\text{sid}} | \tilde{g}_{\text{sid}} | g'_{\text{sid}}$.
- Simulate the random oracle, that is, on input x intended for RO_{384} at the E -interface, if $R(x)$ is not defined then set $R(x) \leftarrow \{0, 1\}^{384}$. Return $R(x)$.
- Deliver messages faithfully for the sessions where the setup is complete, i.e., for clients $C \in \mathcal{C}$ after a successful—i.e., one that was simulated before— $(\text{deliver}, C, \tilde{\eta}, \tilde{m})$ query, and for server sessions $\text{sid} = (\eta, e)$ after delivering a valid group element (valid with respect to the group used in that session) in response to the authenticated message via $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$.

We first observe that the two systems $\mathbf{R} = \prod_{C \in \mathcal{C}} \text{dhec}^C \text{dhes}_{\mathcal{G}}^S [\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}, \text{RO}_{384}]$ and $\mathbf{S} = \sigma^E \text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ are equivalent with respect to how they treat nonces; in particular, the responses to (ack, η) are determined by choosing fresh parameters in both cases. The same argument holds for the queries in which messages are forwarded in the sessions that completed the setup.

For the queries $(\text{deliver}, C, \eta, m)$, for delivering group elements to a session sid of the server, and for querying the random oracle RO_{384} , the difference between \mathbf{R} and \mathbf{S} is that in \mathbf{R} , all “keys” output at either $C \in \mathcal{C}$ or $S/(\eta_C, e)$ for $\eta_C \in \mathcal{N}$ and $e \in [n]$ are chosen consistently with the random oracle queries, whereas in \mathbf{S} they are not (in cases where both group elements have been simulated).

We then describe reduction systems $\mathbf{C}_{i,j}$ for $i \in [|\mathcal{C}|]$ and $j \in [n]$, which obtain at the inside interface a modulus p , a generator g , and two group elements g^a, g^b . We assume that there is some (e.g., lexicographic) ordering on the set \mathcal{C} , i.e., the elements are $C_1, \dots, C_{|\mathcal{C}|}$. All systems $\mathbf{C}_{i,j}$ manage internally a map $R : \{0, 1\}^* \rightarrow \{0, 1\}^{384}$ which is initially empty and is managed

as a “random oracle with lazy evaluation”, i.e., whenever a look-up $R(x)$ for $x \in \{0, 1\}^*$ in R fails, the system will choose a fresh random value $y \in \{0, 1\}^{384}$, store $R(x) = y$, and use y as the result. The system $\mathbf{C}_{i,j}$ then behaves as follows (we stress that every query can be associated with a (client’s) nonce $\eta \in \mathcal{N}$, either because the nonce is given explicitly, or because the query belongs to a session that is described by a nonce and a counter):

- for queries (at the S - and E -interfaces) that are related to nonces $\eta \in \mathcal{N} \setminus \rho(\mathcal{C})$, the system $\mathbf{C}_{i,j}$ can easily reproduce the behavior of the real or ideal systems (their behavior is equivalent for those queries);
- For queries that are related to a client C_ℓ resp. the nonce $\eta_\ell = \rho(C_\ell)$ with $\ell \neq i$, or alternatively to (server) sessions $(\rho(C_i), e)$ with $e \neq j$, emulate the entire sessions. Choose the key output at C_ℓ and the corresponding sub-interface of the server (if there is a consistent session, corresponding to the use of **key-s** in $\text{MSK}_{\mathbf{N}, \rho, \mathcal{F}, \text{SIG}, n}$) by computing as in the protocol and then evaluating R .
- For the j th (**ack**, $\rho(C_i)$)-query, emulate the server’s response using the parameters (i.e., the modulus p , the generator g , and the first group element g^a) obtained at the inside interface.
- For the query (**deliver**, $C_i, \tilde{\eta}, \tilde{m}$) at the sub-interface corresponding to $\succsim \bullet_{\mathbf{N}, \rho, \mathcal{F}, \text{SIG}, n}$:
 - if $\tilde{\eta} = \eta_{sid}$ for $sid = (\rho(C_i), j)$, and if the message would be admitted by $\succsim \bullet_{\mathbf{N}, \rho, \mathcal{F}, \text{SIG}, n}$, then emulate the second group element g^b obtained at the inside interface as a message from C_i to sid . Choose $\kappa_{C_i} \in_R \{0, 1\}^{384}$ and emulate it at interface C_i .
 - if $\tilde{\eta} = \eta_{sid}$ for $sid = (\rho(C_i), e)$ and $e \neq j$, and if the message would be admitted by $\succsim \bullet_{\mathbf{N}, \rho, \mathcal{F}, \text{SIG}, n}$, then emulate a uniformly chosen element of the group described in \tilde{m} ; choose the key κ_{C_i} using the map R evaluated on $\tilde{g}|\text{master secret}|\rho(C_i)|\eta_{sid}$, where \tilde{g} is the DH key that can be computed because $\mathbf{C}_{i,j}$ chose all parameters.
- Upon delivery of the first client’s message in session $sid = (\rho(C_i), j)$, if the message completes a consistent session with C_i , then emulate the output κ_{C_i} at the sid -interface. Otherwise, compute the server’s output according to the protocol.
- Queries $x \in \{0, 1\}^*$ to RO_{384} are answered by evaluating $y = R(x)$ as described above and answering with y . If $x = \tilde{g}|\text{master secret}|\rho(C_i)|\eta_{sid}$, where \tilde{g} is a valid group element with respect to the group \mathbf{Z}_p , use the queries at the inside interface to determine whether $g^{ab} = \tilde{g}$ and, in case of success, input \tilde{g} as solution at the inside interface and halt.

In the following, we use $\mathbf{C}_{*,*}$ wherever the indices i, j of $\mathbf{C}_{i,j}$ are irrelevant. Consider now the following MBOs $\mathcal{E}^{i,j} = (E_1^{i,j}, E_2^{i,j}, \dots)$ for each pair (i, j) with $i \in [|\mathcal{C}|]$ and $j \in [n]$, such that $E_q^{i,j}$ is defined over the random systems $\mathbf{C}_{*,*} \left(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}} \right)^{-}$ (and \mathbf{R}, \mathbf{S}):¹¹ The MBO becomes 1 if after q queries there has been a query ($\text{pmk}|\text{master secret}|\eta_C|\eta_{sid}$) at the random oracle, such that $C = C_i$, $sid = (\eta_C, j)$ and pmk is the DH element corresponding to session (i, j) , i.e., the DH element with respect to the server’s group element sent in session (η_C, j) and the group element sent by C_i (given they are in the same group). Then define the MBO \mathcal{E} via $E_q = \bigvee_{i,j} E_q^{i,j}$.

By definition of $\mathbf{C}_{i,j}$ and \mathcal{E} , one can verify that $\mathbf{R}^{\mathcal{E}} \stackrel{g}{\equiv} \mathbf{C}_{i,j}^{\mathcal{E}} \left(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}} \right)^{-} \stackrel{g}{\equiv} \mathbf{S}^{\mathcal{E}}$ for all pairs $(i, j) \in [|\mathcal{C}|] \times [n]$. Also, provoking the MBO $\mathcal{E}^{i,j}$ in $\mathbf{C}_{i,j} \left(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}} \right)^{-}$ implies that the reduction

¹¹The term $\left(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}} \right)^{-}$ refers to the game $\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}$ without its Monotone Binary Output (MBO), see the last paragraph of the proof.

$\mathbf{C}_{i,j}$ is successful in winning $\mathbf{G}_G^{\text{GapDH}}$, so

$$\Gamma^{\mathbf{DC}_{i,j}}(\mathbf{G}_G^{\text{GapDH}}) = \Gamma^{\mathbf{D}}(\mathbf{C}_{i,j}\mathbf{G}_G^{\text{GapDH}}) \geq \Gamma^{\mathbf{D}}\left(\mathbf{C}_{i,j}^{\mathcal{E}^{i,j}}\left(\mathbf{G}_G^{\text{GapDH}}\right)^{-}\right)$$

and

$$\sum_{i,j} \Gamma^{\mathbf{D}}\left(\mathbf{C}_{i,j}^{\mathcal{E}^{i,j}}\left(\mathbf{G}_G^{\text{GapDH}}\right)^{-}\right) \geq \Gamma^{\mathbf{D}}\left(\mathbf{C}_{*,*}^{\mathcal{E}}\left(\mathbf{G}_G^{\text{GapDH}}\right)^{-}\right)$$

by Lemma 27. The statement then follows using Lemma 26 and using the reduction \mathbf{C} that chooses any one of the $\mathbf{C}_{i,j}$ with $i \in [|\mathcal{C}|]$ and $j \in [n]$ uniformly at random. \square

3.3.3 The RSA-PKCS Scheme

The RSA-PKCS-based protocol consists of the following two converters, which are based on the RSA algorithms $\text{RSA} = (\text{gen}, \text{enc}, \text{dec})$ as specified in PKCS#7 [Kal98] (see also Appendix A.2). The client's converter **rsac** behaves as follows:

1. Obtain the nonces η_C and η_{sid} from $\text{SNET}_{\mathbf{N},\rho,n}$.
2. Obtain a message $\text{cert} = (pk, s)$ (supposed to be the server's certificate) via $\text{SNET}_{\mathbf{N},\rho,n}$.
3. Verify the server's certificate by querying $(\text{verify}, \text{cert})$ at $\text{PKI}_{\mathfrak{F}}$ (if that fails, abort).
4. Choose a secret $\text{pmk} \in \{0,1\}^{384}$ as follows: concatenate the two-byte protocol version identifier with a 46-byte uniformly random string, encrypt with the server's public key pk obtained from cert (resulting in a ciphertext $c = \text{enc}(\text{pmk}; pk)$), and send c via $\text{SNET}_{\mathbf{N},\rho,n}$.
5. Query $\text{pmk}|\text{master secret}|\eta_C|\eta_{sid}$ at RO_{384} , call the result κ .
6. Output $(\kappa, \eta_C, \eta_{sid}, (\text{cert}, c))$ at the outside interface.
7. Forward all further communication between the outside interface and $\text{SNET}_{\mathbf{N},\rho,n}$.

The server's converter **rsas** behaves as follows:

0. Generate an RSA key pair $(pk, sk) = \text{gen}$ and register the public key pk at $\text{PKI}_{\mathfrak{F}}$ via $(\text{register}, pk)$, obtaining a value s , define $\text{cert} = (pk, s)$.
- For each of the sessions described by a pair $\text{sid} = (\eta_C, e) \in \mathcal{N} \times [n]$:
1. Upon receiving a nonce η_{sid} at sub-interface $\text{sid} = (\eta_C, e)$, respond with the certificate cert .
 2. Upon receiving a ciphertext c , decrypt $\tilde{\text{pmk}} = \text{dec}(c; sk)$. If $\tilde{\text{pmk}} = \perp$ or if the first two bytes of $\tilde{\text{pmk}}$ do not match the two-byte protocol version identifier, choose a fresh value for $\tilde{\text{pmk}}$ as in Step 4. of **rsac**.
 3. Query $\tilde{\text{pmk}}|\text{master secret}|\eta_C|\eta_{sid}$ at RO_{384} and call the result $\tilde{\kappa}$.
 4. Output $(\tilde{\kappa}, \eta_{sid}, (\text{cert}, c))$ at the $(\text{sid}$ -sub-interface of the) outside interface.
 5. Forward all further communication between the $(\text{sid}$ -sub-interface of the) outside interface and (the sid -sub-interface of) $\text{SNET}_{\mathbf{N},\rho,n}$.

We show that the described protocol is secure under the NR-PCA assumption for PKCS#7, i.e., it constructs the resource $\text{MSK}_{\mathbf{N},\rho,AUX,n}$. Here, the distribution AUX in this case consists of two parts. The first part is the same for all sessions and consists of the certificate (i.e., depends on \mathfrak{F} of $\text{PKI}_{\mathfrak{F}}$). The distribution is described by $(sk, pk) \leftarrow \text{gen}$, $f \leftarrow \mathfrak{F}$, and then $\text{cert} = (pk, f(pk))$. The second part is chosen independently for each session by doing the same process as in the protocol: choose $\text{pmk} \in \{0,1\}^{384}$ (as it is done in Step 4. of **rsac**) and compute $c = \text{enc}(\text{pmk}; pk)$.

Lemma 12. *The protocol $(\text{rsac}, \text{rsas})$ constructs from $[\text{RO}_{384}, \text{SNET}_{\mathbf{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}]$ the master secret key resource $\text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$, under the NR-PCA assumption for $\text{RSA} = (\text{gen}, \text{enc}, \text{dec})$, i.e., for the simulator σ and the reductions \mathbf{C}_q described in the proof,*

$$[\text{RO}_{384}, \text{SNET}_{\mathbf{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}] \xrightarrow{(\text{rsac}, \text{rsas}), \sigma, (0, \varepsilon)} \text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n},$$

where for each distinguisher \mathbf{D} , $\varepsilon(\mathbf{D}) = n \cdot |\mathcal{C}| \cdot \Gamma_q^{\text{DC}_q}(\mathbf{G}^{\text{nr-pca}}) + \frac{q}{2^{368}}$.

Before proceeding with the formal proof, let us discuss some intuition. In the proofs of Lemma 9 and Lemma 11, it is apparent when the pre-master secret value embedded in the challenge to the reduction is delivered to the server. However, as we argue below, this is not necessarily the case for TLS-RSA if one assumes like [JK02, KPW13] that RSA is OW-PCA. Recall that in the RSA mode, the client chooses a random pre-master secret and encrypts it under the server's certified public key. Now, the argument would be that if the distinguisher can distinguish between the $\text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ resource and the construction based on the RSA encryption, then it can break the OW-PCA encryption (essentially it knows the pre-master secret). The reduction to OW-PCA uses a “testing” $\mathbf{PCA}(\cdot, \cdot)$ oracle, such that $\mathbf{PCA}(m, c)$ outputs 1 iff. $\text{dec}(c; sk) = m$; in our case, this should permit the reduction to properly simulate the random oracle.

In the case of RSA, there is a subtle problem which does not appear in the DH case. In particular, whereas DH elements are not re-randomizable, if the OW-PCA encryption scheme is re-randomizable, the distinguisher between two subsequent hybrids can inject a re-randomized ciphertext which cannot be simulated accurately by the reduction. Recall that in the DHE case, the hybrids were indexed by sessions, such that the input from the GapDH game was inserted into a different session each time. If the distinguisher receives a challenge ciphertext from the OW-PCA reduction, then re-randomizes it and sends it to the server in the “correct session,” then the behavior of the two resources will differ, since for one of them the key output of the client and server will coincide, whereas for the other, it will differ (recall that since the randomized ciphertext will decrypt to the same plaintext as the challenge ciphertext, the pre-master secret input into the random oracle is the same). For this reason, similar to [BFK⁺13b], Lemma 12 above relies on the assumption that RSA-PKCS is not re-randomizable.¹²

Proof.

Availability. We first prove the availability condition. The “ideal” system $\perp^E \text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ chooses for each session $\text{sid} = (\rho(C), 1)$ a nonce η_{sid} , a key κ_C and auxiliary information aux_C , and outputs $(\kappa_C, \rho(C), \eta_{\text{sid}}, \text{aux}_C)$ at the C interface and $(\kappa_C, \eta_{\text{sid}}, \text{aux}_C)$ at the interface S/sid ; afterwards the communication is forwarded between C and S/sid .

The distribution in the case $\prod_{C \in \mathcal{C}} \text{rsac}^C \text{rsas}^S [\text{RO}_{384}, \text{SNET}_{\mathbf{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}]$ is as follows. The key κ_C for each client C is a uniformly random 384-bit string (the random oracle RO_{384} is queried at distinct places as $\eta_C = \rho(C)$ is distinct by the assumption on ρ), the nonces η_{sid} are chosen according to the distribution \mathbf{N} , and the auxiliary information aux_C has exactly the same distribution as well. Moreover, after at both interfaces C and $S/(\eta_C, 1)$ the above information is output, the resource behaves as a bidirectional channel between those interfaces. This verifies the validity of the availability condition.

¹²An alternative approach would be to be less modular and consider the entire protocol as a single unit, since only the finished messages authenticate the exact ciphertext transmitted during the session. This is essentially the path taken by Krawczyk *et al.* [KPW13].

Security. In order to prove the validity of the security condition, consider the following simulator $\sigma = \sigma(\rho)$ (parametrized by the function ρ of $\text{MSK}_{\mathbf{N},\rho,AUX,n}$):

- Initially, σ sets $e_\eta = 1$ for all $\eta \in \mathcal{N}$ and defines an (initially empty) map $R : \{0,1\}^* \rightarrow \{0,1\}^{384}$.
- Choose a function $f \leftarrow \mathfrak{F}$, sample $(pk, sk) \leftarrow \text{gen}$, compute $s = f(pk)$ and output $\text{cert} = (pk, s)$ at the outside interface to simulate the answer of an input $(\text{register}, pk)$ to $\text{PKI}_{\mathfrak{F}}$.
- Upon input (ack, η) for a nonce $\eta \in \mathcal{N}$ at the outside interface, set $\text{sid} = (\eta, e_\eta)$. If $e_\eta \leq n$, then issue (ack, η) at the inside interface. Obtain the nonce η_{sid} at the inside interface and output η_{sid} as response of $\text{SNET}_{\mathbf{N},\rho,n}$ at the outside interface; increment e_η . Also, output cert at the outside interface as being transmitted via $\text{SNET}_{\mathbf{N},\rho,n}$ from sid .
- Upon input $(\text{deliver}, C, \tilde{\eta})$ and after cert has been delivered via $\text{SNET}_{\mathbf{N},\rho,n}$ at the outside interface, issue $(\text{key-c}, C, \text{aux}, \tilde{\eta})$ at the inside interface. The auxiliary information aux consists of $\text{aux} = (\text{cert}, c)$ where the ciphertext c is computed by encrypting a value pmk (as in Step 4. of rsac) using public key pk . Output c at the outside interface as being transmitted via $\text{SNET}_{\mathbf{N},\rho,n}$ from C and register $(C, \tilde{\eta}, \text{pmk})$.
- When obtaining a ciphertext \tilde{c} at the outside sub-interface sid , with $\text{sid} = (\eta, e)$, compute $\tilde{\text{pmk}} = \text{dec}(\tilde{c}; sk)$. If $\text{dec}(\tilde{c}; sk) = \perp$ or if the first two bytes of $\tilde{\text{pmk}}$ do not match the two-byte protocol version identifier, then choose $\tilde{\text{pmk}}$ uniformly at random instead. Then:
 - If $\eta \in \rho(C)$ and with $C = \rho^{-1}(\eta)$, in case $\tilde{\text{pmk}}$ corresponds to a previously recorded triple $(C, \eta_{\text{sid}}, \tilde{\text{pmk}})$, input $(\text{key-s}, \eta, e, \text{aux})$ at the inside interface, with aux being (cert, \tilde{c}) .
 - Otherwise, input $(\text{inject}, \eta, e, \text{aux}, \kappa)$ at the inside interface, with $\text{aux} = \text{cert}|\tilde{c}$ with the ciphertext \tilde{c} just injected. The value κ is defined as $R(\tilde{\text{pmk}}|\text{master secret}|\eta|\eta_{\text{sid}})$. (If $\tilde{\text{pmk}}|\text{master secret}|\eta|\eta_{\text{sid}}$ has not been queried to the random oracle before, define it as $R(\tilde{\text{pmk}}|\text{master secret}|\eta|\eta_{\text{sid}}) \leftarrow_{\$} \{0,1\}^{384}$.)
- Simulate the random oracle, that is, on input x intended for RO_{384} at the E -interface, if $R(x)$ is not defined then set $R(x) \leftarrow_{\$} \{0,1\}^{384}$. Return $R(x)$.

Define $\mathbf{R} \equiv \prod_{C \in \mathcal{C}} \text{rsac}^C \text{rsas}^S[\text{RO}_{384}, \text{SNET}_{\mathbf{N},\rho,n}, \text{PKI}_{\mathfrak{F}}]$ and $\mathbf{S} \equiv \sigma^E \text{MSK}_{\mathbf{N},\rho,AUX,n}$. We start by noting that the behavior of the two systems \mathbf{R} and \mathbf{S} is identical for the initial queries:

- The value cert has the same distribution in both resources.
- The responses to (ack, η_C) are determined by choosing fresh parameters in both cases.
- The distribution corresponding to messages forwarded in the sessions that completed the setup is also the same in the two systems.

For the queries $(\text{deliver}, C, \tilde{\eta})$, for delivering a ciphertext to a session sid of the server, and for querying the random oracle RO_{384} , the difference between \mathbf{R} and \mathbf{S} is that in \mathbf{R} all “keys” output at either $C \in \mathcal{C}$ or $S/(\eta, e)$ for $\eta \in \mathcal{N}$ and $e \in [n]$ are chosen consistently with the random oracle queries, whereas in \mathbf{S} they are not (in cases where the ciphertext has been simulated). The intuition behind the proof is that a distinguisher telling apart the two systems would have to query the random oracle at such a point; however, as we argue below, the latter is not very likely to happen as otherwise such a distinguisher can be used to break the NR-PCA assumption for $\text{RSA} = (\text{gen}, \text{enc}, \text{dec})$ (see Appendix C.1). Recall that in this game the challenger produces

a ciphertext c^* corresponding to the encryption of a random message m^* using public key pk , and the goal of the adversary is to produce a re-randomized ciphertext $c' \neq c^*$ decrypting to m^* , with the help of a $\mathbf{PCA}(\cdot, \cdot)$ oracle (where $\mathbf{PCA}(m, c)$ outputs 1 iff $\text{dec}(c; sk) = m$).¹³

We now prove the security condition. We describe a series of reductions $\mathbf{C}_{i,j,q}$ with $i \in [\mathcal{C}]$, $j \in [n]$, and $q \in \mathbb{N}$, where the indices pinpoint a session run for a nonce η corresponding to a client $C \in \mathcal{C}$ and a corresponding nonce η_{sid} generated for session $sid = (\eta, e)$, where $1 \leq e \leq e_\eta \leq n$. We assume some (e.g. lexicographic) order \preceq over the set of clients, and write C_i to denote the i -th client with respect to this order. The reduction $\mathbf{C}_{i,j,q}$ uses the public key pk and injects the challenge c^* received from $\mathbf{G}^{\text{nr-pca}}$, relies on the PCA queries it can make at $\mathbf{G}^{\text{nr-pca}}$, and it works as follows:

- Let (pk, c^*) be the input received from the inside interface with $\mathbf{G}^{\text{nr-pca}}$.
- For all the sessions between a client $C \neq C_i$ and for the sessions of C_i with server interfaces (η, e) such that $e \neq j$, $\mathbf{C}_{i,j,q}$ uses pk to encrypt a pmk and generates the clients' keys as in the real resource.
- In the session between client C_i and the server session (η, j) , the session is emulated using c^* . In this case the output at the client and server interface is independent of the queries made to RO_{384} .

The simulation of the random oracle queries and server master secret keys goes into more details as follows. Each of the reductions will keep two (initially empty) lists \mathcal{H}_P and \mathcal{H}_C . List \mathcal{H}_P contains entries of the form $(pmk, \eta_C, \eta_{sid}, y)$ while \mathcal{H}_C contains entries of the form $(c, \eta'_C, \eta'_{sid}, y')$. \mathcal{H}_P corresponds to random oracle queries, while \mathcal{H}_C corresponds to master secret keys output at the server. The two lists need to be kept consistent, in the sense that if for two entries $\eta_C = \eta'_C$, $\eta_{sid} = \eta'_{sid}$, and $\mathbf{PCA}(pmk, c) = 1$ then $y = y'$. Hence:

- Whenever the random oracle is queried on input $pmk|\text{master secret}|\eta_C|\eta_{sid}$ at RO_{384} , first look for an entry $(pmk, \eta_C, \eta_{sid}, y)$ in \mathcal{H}_P ; return the corresponding value y in case such entry is found. Check for the entry $(\tilde{c}, \eta_C, \eta_{sid}, y)$ in \mathcal{H}_C , and return the corresponding y if $\mathbf{PCA}(pmk, \tilde{c}) = 1$.

Otherwise, check if $\mathbf{PCA}(pmk, c^*) = 1$:

- If this is the case, run $c' \leftarrow \text{enc}(pmk; pk)$ with different randomness until $c' \neq c^*$ and output c' as the re-randomized ciphertext in $\mathbf{G}^{\text{nr-pca}}$; note that, according to PKCS#7, the above re-sampling requires to encrypt the plaintext with a different random pad P' and two trials are sufficient.
- Else, respond with a random y and update the list \mathcal{H}_P accordingly.
- When some ciphertext \tilde{c} is delivered to a server's session $sid = (\eta, e)$, search for an entry $(\tilde{c}, \eta, \eta_{sid}, y)$ in \mathcal{H}_C . Otherwise, check whether there is an entry $(pmk, \eta, \eta_{sid}, y)$ in \mathcal{H}_P such that $\mathbf{PCA}(pmk, \tilde{c}) = 1$. Return the corresponding value y in case either of these entries is found; else, respond with a random y and update the list \mathcal{H}_C accordingly.
- If after q queries no ciphertext c' has been returned to $\mathbf{G}^{\text{nr-pca}}$, check whether \mathcal{H}_C contains an entry $(c, \eta_C, \eta_{sid}, *)$ such that $C = C_i$, $sid = (\eta_C, j)$ and return c to $\mathbf{G}^{\text{nr-pca}}$.

As for DHE, we use $\mathbf{C}_{*,*,q}$ wherever the indices i, j of $\mathbf{C}_{i,j,q}$ are irrelevant and we consider a monotone binary output (MBO) $\mathcal{E}^{i,j} = (E_1^{i,j}, E_2^{i,j}, \dots)$ for each pair (i, j) with $i \in [\mathcal{C}]$ and $j \in [n]$. Here, $E_q^{i,j}$ is defined over the random systems $\mathbf{C}_{*,*,q'}(\mathbf{G}^{\text{nr-pca}})^-$ (and \mathbf{R}, \mathbf{S}):¹⁴ The

¹³Note that in TLS the message space is of the form $\mathcal{M} = \text{version_number} \times \{0, 1\}^{368}$. The NR-PCA assumption over this space is implied by the one on $\{0, 1\}^{386}$, but one would loose an additional term $\frac{1}{2^{16}}$ in the reduction.

¹⁴Note that q' ranges over $1, 2, \dots$ independently of q as the MESs are defined for all reductions.

MBO becomes 1 if after q queries there has been a query $(pmk|master\ secret|\eta_C|\eta_{sid})$ at the random oracle, such that $C = C_i$, $sid = (\eta_C, j)$ and $dec(c^*, sk) = pmk$ or a ciphertext c has been delivered in session $sid = (\eta_C, j)$ with $C = C_i$ such that $dec(c^*, sk) = dec(c; sk)$.¹⁵

We also define three MBOs $\mathcal{F} = (F_1, F_2, \dots)$, $\mathcal{F}' = (F'_1, F'_2, \dots)$, and $\mathcal{F}'' = (F''_1, F''_2, \dots)$ which we need to “unify” slightly different distributions of outputs with respect to injections of invalid ciphertexts. A ciphertext is invalid if it decrypts to \perp or if the first two bytes of its plaintext do not match the two-byte protocol version identifier. In both **R** and **S**, if an invalid ciphertext is injected in a session with nonces η_C and η_{sid} , a random pre-master secret pmk^* is chosen and the random oracle (either RO_{384} or R) is evaluated on the value $pmk^*|master\ secret|\eta_C|\eta_{sid}$. Note that if the pre-master secret pmk^* collides with a value that was queried in combination with the same nonces η_C and η_{sid} before (either because it was computed by the client or because it was queried directly at RO_{384} by the attacker), the output of the random oracle, and hence at the server’s interface, would of course be consistent with the previously obtained value. If the value pmk^* has not been used before, the response of the random oracle and hence the key output at the server’s interface are uniformly random and independent of all previous values. The reduction $\mathbf{C}_{*,*,*}$, however, cannot distinguish invalid ciphertexts from ciphertexts decrypting to a pre-master secret that has not been used before (the PCA query to \mathbf{G}^{nr-pca} returns the same value on both), and hence always outputs a uniformly random key.

The pre-master secret pmk^* has 368 bits of entropy. Thus, after r different pre-master secrets have been queried at the random oracle with the same nonces, the probability for the freshly chosen pmk^* to collide with a previously used value is $r \cdot 2^{-368}$. With the remaining probability $1 - r \cdot 2^{-368}$, the query to the random oracle is fresh and the output distribution is uniform (i.e., every value has probability 2^{-384}). In case of $\mathbf{C}_{*,*,*}$, an invalid ciphertext always leads to a uniformly random output. As a result, an output of **R** or **S** on such a “dangerous” query with respect to an invalid ciphertext is slightly more likely to collide with previous outputs.

We rectify the above difference in probabilities by defining the three monotone binary outputs \mathcal{F} , \mathcal{F}' , and \mathcal{F}'' (for **R**, **S**, and $\mathbf{C}_{*,*,*}$, respectively) as follows.

- The MBO \mathcal{F} for **R** becomes 1 once in a session with nonces η_C and η_{sid} , an invalid ciphertext \tilde{c} was input and lead to a query $pmk^*|master\ secret|\eta_C|\eta_{sid}$ (where pmk^* was chosen uniformly at random) to RO_{384} , and the same query was also asked at RO_{384} either by $rsac$ or via the E -interface.
- The MBO \mathcal{F}' for **S** becomes 1 once in a session with nonces η_C and η_{sid} , an invalid ciphertext \tilde{c} was input and lead to an evaluation on $pmk^*|master\ secret|\eta_C|\eta_{sid}$ (where pmk^* was chosen uniformly at random) of R within σ , and the simulator σ computed an RSA ciphertext from $C = \rho^{-1}(\eta_C)$ (if defined) with the same pre-master secret pmk^* or R was also evaluated on the same value because of a query at the outside interface.

The above two MBO \mathcal{F} (resp. \mathcal{F}') guarantee that, given that the MBO remains 0 during a dangerous query, the output distribution is uniform. Using the value r defined above, this means that, for every possible key κ , the probability of the MBO remaining 0 and the key κ appearing is $(1 - r \cdot 2^{-368}) \cdot 2^{-384}$. What remains to be done is defining an MBO \mathcal{F}'' on $\mathbf{C}_{*,*,*}$ that leads to the same output distribution for such queries. But as the distribution in that case is uniform anyways, all that remains to be done is computing the probability for the MBO to be provoked during a “dangerous” query (see below), and provoking the MBO with the respective probability (independently of the value output by $\mathbf{C}_{*,*,*}$. The probability is computed in more detail as follows:

¹⁵Note that the event is well defined since for the RSA-PKCS scheme sk is uniquely defined by pk .

- if the query is the first query $pmk'|\text{master secret}|\eta|\tilde{\eta}$ for this value of pmk' to RO_{384} , where the first two bytes of pmk' match the two-byte protocol version identifier and an invalid ciphertext was sent to session $sid = (\eta, e)$ with nonce $\eta_{sid} = \tilde{\eta}$ before, then \mathcal{F}''' becomes 1 with probability 2^{-368} ;
- if the query is sending an invalid ciphertext c in session (η, e) , then the probability is computed by first counting the number r of *distinct* queries to RO_{384} and (potentially) the output at the corresponding client $\rho^{-1}(\eta)$ that collide in terms of generated keys; then \mathcal{F} becomes 1 with probability $r \cdot 2^{-368}$.

Then define the MBOs \mathcal{E} via $E_q = \bigvee_{i,j} E_q^{i,j} \vee F_q$, \mathcal{E}' via $E'_q = \bigvee_{i,j} E_q^{i,j} \vee F'_q$, and \mathcal{E}'' via $E''_q = \bigvee_{i,j} E_q^{i,j} \vee F''_q$.

By the definition of $\mathbf{C}_{i,j,q}$, \mathcal{E} , \mathcal{E}' , and \mathcal{E}'' , one can verify that $\mathbf{R}^{\mathcal{E}} \stackrel{g}{=} \mathbf{C}_{i,j,q}^{\mathcal{E}''} (\mathbf{G}^{\text{nr-pca}})^- \stackrel{g}{=} \mathbf{S}^{\mathcal{E}'}$ for all pairs $(i, j, q) \in [|\mathcal{C}|] \times [n] \times \mathbb{N}$. Also, provoking the MBO $E_q^{i,j}$ in $\mathbf{C}_{i,j,q} (\mathbf{G}^{\text{nr-pca}})^-$ implies that the reduction $\mathbf{C}_{i,j}$ is successful in winning $\mathbf{G}^{\text{nr-pca}}$, so

$$\Gamma_q^{\mathbf{D}\mathbf{C}_{i,j,q} (\mathbf{G}^{\text{nr-pca}})} = \Gamma_q^{\mathbf{D}} (\mathbf{C}_{i,j,q} \mathbf{G}^{\text{nr-pca}}) \geq \Gamma_q^{\mathbf{D}} (\mathbf{C}_{i,j,q}^{\mathcal{E}''} (\mathbf{G}^{\text{nr-pca}})^-)$$

and $\sum_{i,j} \Gamma_q^{\mathbf{D}} (\mathbf{C}_{i,j,q}^{\mathcal{E}''} (\mathbf{G}^{\text{nr-pca}})^-) + \Gamma_q^{\mathbf{D}} (\mathbf{C}_{i,j,q}^{\mathcal{F}''} (\mathbf{G}^{\text{nr-pca}})^-) \geq \Gamma_q^{\mathbf{D}} (\mathbf{C}_{*,*,q}^{\mathcal{E}''} (\mathbf{G}^{\text{nr-pca}})^-)$ by Lemma 27. The statement then follows using Lemma 26 and using the reduction \mathbf{C}_q that chooses any one of the $\mathbf{C}_{i,j,q}$ with $i \in [|\mathcal{C}|]$ and $j \in [n]$ uniformly at random. The probability of provoking \mathcal{F} , \mathcal{F}' , resp. \mathcal{F}'' can be upper bounded by $q \cdot 2^{-368}$. \square

3.3.4 The Handshake in TLS 1.3

In this section, we consider a slightly modified version of the handshake in the suggested TLS 1.3. What we modify is that the server's **Certificate** and **CertificateVerify** are not encrypted. This does not seem to have any impact on practical security beyond the statements proven here (the server's certificate is public anyway, and the **CertificateVerify** message is basically a signature on a hash of all previous messages), but it allows us to make a security proof.

The modified authenticated transmission resource. As in Section 3.3.2, we begin by modeling the construction that is achieved by the signature scheme, namely the construction of a resource $\triangleright \bullet'$ that allows the server to transmit, in each session, a single message to the client authentically. Unfortunately, we cannot re-use the resource and construction shown in Section 3.3.2, for two reasons: first, since the signature computation now also includes the first client message, this message is authenticated as well (this is required in the subsequent construction step); second, the computation of the signature is changed. In TLS 1.2, only the “contents” of the previously exchanged fragments is authenticated (but not, e.g., the packet headers). In contrast, TLS 1.3 authenticates the complete fragments, including all headers.

The resource $\triangleright \bullet'$ has (client) interfaces $C \in \mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$, a (server) interface S with sub-interfaces $(\eta, e) \in \mathcal{N} \times \mathbb{N}$ with $e \in \mathbb{N}$, and an (eavesdropper) interface E with sub-interfaces $\mathcal{A}_{\text{TCP}} \cup (\mathcal{N} \times \mathbb{N})$, and is parametrized by an injection $\rho : \mathcal{C} \rightarrow \mathcal{N}$, a distribution \mathfrak{F} over functions $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, a signature scheme $SIG = (gen, sign, vrf)$ as defined in Appendix B.1, and a hash function $h : \{0, 1\}^* \rightarrow \{0, 1\}^k$. The resource is described in detail in Figure 6.

The following protocol constructs $\triangleright \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n, h}$ from $\mathbf{SNET}_{\mathbf{N}, \rho, n}$ and $\mathbf{PKI}_{\mathfrak{F}}$: The client's converter vrf is based on the signature scheme SIG and behaves as follows:

0. Obtaining the nonce $\eta_C \in \mathcal{N}$ from $\mathbf{SNET}_{\mathbf{N}, \rho, n}$, output η_C at the outside interface. Forward all messages m_1, m_2, \dots from the outside to the inside interface. (Messages sent from the client to the server.)

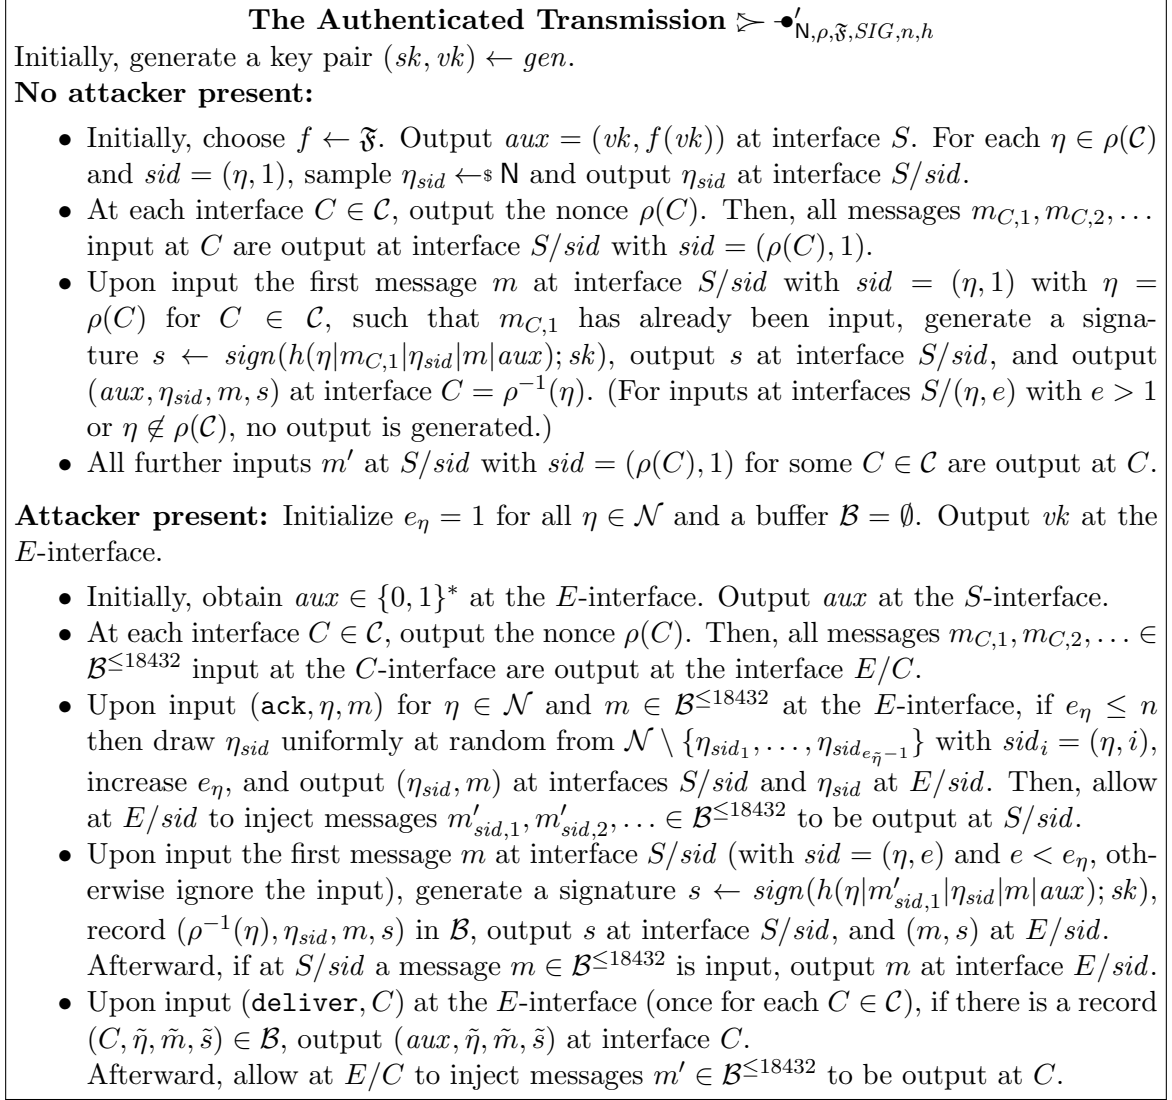


Figure 6: The network allowing the server to transmit, to each client, one message authentically. TLS 1.3 variant.

1. Obtain the server's nonce $\eta_{sid} \in \mathcal{N}$ from $\text{SNET}_{\mathcal{N}, \rho, n}$.
2. Obtain the first message (m, cert, s) from $\text{SNET}_{\mathcal{N}, \rho, n}$. Query $(\text{verify}, \text{cert})$ at $\text{PKI}_{\mathfrak{F}}$; abort if the verification fails or if cert is not a well-formed certificate $\text{cert} = (vk, f(vk))$. If $\text{vrf}(h(\eta_C|m_1|\eta_{sid}|m|\text{cert}), s; vk) = 1$, then output $(\text{cert}, \eta_C, \eta_{sid}, m, s)$ at the outside interface. (Otherwise abort.)
3. Forward all further messages from the inside to the outside interface. (Messages sent from the server to the client.)

The server's converter **sgn** provides at the outside "sessions" for all $sid = (\eta_C, e) \in \mathcal{N} \times [n]$ and behaves as follows:

0. Compute $(sk, vk) \leftarrow \text{gen}$. Input vk at $\text{PKI}_{\mathfrak{F}}$, obtaining a response s , and set $\text{cert} = (vk, s)$. Output cert at the outside interface (as auxiliary information).
- For each session $sid = (\eta_C, e)$ —i.e., the inputs/outputs at $\text{SNET}_{\mathcal{N}, \rho, n}$ are at the corresponding inside sub-interface sid , and the inputs/outputs at the outside interface are at sub-interface

sid —do:

1. Receiving a nonce η_{sid} and the first message m from $\text{SNET}_{\mathbf{N},\rho,n}$, output (η_{sid}, m) at the outside.
2. Forward all further messages $m'_{sid,1}, m'_{sid,2}, \dots$ from the outside interface to the inside interface.
3. Obtaining the first input m at the outside, compute $s \leftarrow \text{sign}(h(\eta_C | m'_{sid,1} | \eta_{sid} | m | \text{cert}); sk)$ and send (m, cert, s) via $\text{SNET}_{\mathbf{N},\rho,n}$. Output s at the outside.
4. Forward further messages from the outside interface to the inside interface.

Lemma 13. *The protocol (vrf, sgn) for a particular signature scheme $\text{SIG} = (\text{gen}, \text{sign}, \text{vrf})$ constructs $\succsim \bullet_{\mathbf{N},\rho,\mathfrak{F},\text{SIG},n}$ from $\text{SNET}_{\mathbf{N},\rho,n}$ and $\text{PKI}_{\mathfrak{F}}$, if the signature scheme SIG is unforgeable and the hash function h is collision resistant. In more detail, for the simulator σ and the reductions \mathbf{C} and \mathbf{C}' described in the proof,*

$$[\text{SNET}_{\mathbf{N},\rho,n}, \text{PKI}_{\mathfrak{F}}] \xrightarrow{(\text{vrf}, \text{sgn}), \sigma, (0, \varepsilon)} \succsim \bullet'_{\mathbf{N},\rho,\mathfrak{F},\text{SIG},n,h},$$

with $\varepsilon(\mathbf{D}) \doteq \Gamma^{\mathbf{DC}}(\mathbf{G}^{\text{uf-cma}}) + \Gamma^{\mathbf{DC}'}(\mathbf{G}^{\text{CR}})$ for all distinguishers \mathbf{D} .

Proof.

Availability. The availability condition follows almost as in Lemma 10. The only differences are that the first message in each session must always be given by the client $C \in \mathcal{C}$, and that the signatures are computed differently.

Security. For proving the security condition, we consider the following simulator σ :

- Initially, σ obtains the public key vk at the inside interface, chooses a function $f \leftarrow \mathfrak{F}$, and computes $\text{cert} = (vk, f(vk))$. Also, σ internally initializes $e_\eta = 1$ for all $\eta \in \mathcal{N}$, and inputs cert as auxiliary information at the inside interface, and as output of $\text{PKI}_{\mathfrak{F}}$ at the outside interface.
- For clients $C \in \mathcal{C}$, messages $m_{C,1}, m_{C,2}, \dots \in \mathcal{M}$ obtained at sub-interface C at the inside are output at the outside as being sent by C via $\text{SNET}_{\mathbf{N},\rho,n}$.
- Upon input (ack, η) for a nonce $\eta \in \mathcal{N}$ at the outside interface, set $sid = (\eta, e_\eta)$. If $e_\eta \leq n$, then, upon the first message m' directed to the server session sid obtained at the outside interface, input (ack, η, m') at the inside interface. Obtain the nonce η_{sid} at the inside interface, output η_{sid} as response of $\text{SNET}_{\mathbf{N},\rho,n}$ at the outside interface, and increase e_η . Then, further messages $m'_{sid,1}, m'_{sid,2}, \dots$ directed to the server session sid obtained at the outside interface are also output at the inside.
- Upon input $(\text{deliver}, C, \tilde{\eta})$ at the outside sub-interface corresponding to $\text{SNET}_{\mathbf{N},\rho,n}$, if this is the first such query for C , record the nonce as $\eta_{sid} \doteq \tilde{\eta}$ for $sid = (\rho(\eta), *)$.
- When obtaining the first output (m, s) at the inside sub-interface sid with $sid = (\eta, e)$, output (m, cert, s) as being transmitted on $\text{SNET}_{\mathbf{N},\rho,n}$ as the second message via the corresponding sub-interface. Mark sid as “active”.
- When receiving at the outside sub-interface corresponding to $\text{SNET}_{\mathbf{N},\rho,n}$ the first message $\tilde{m}_1 = (m', \text{cert}', s')$ to a client C , if $\text{cert}' = \text{cert}$ and, with $sid = (\eta_C, *)$, the verification $\text{vrf}(h(\eta_C | m'_{sid,1} | \eta_{sid} | m' | \text{cert}), s'; vk) = 1$ succeeds, then input $(\text{deliver}, C)$ at the inside interface¹⁶ and mark C as “active”.

¹⁶This has an effect only if the corresponding message has been sent by the server before; if the signature is forged, the real and ideal systems behave differently.

- Further server-to-client communication is then also forwarded: Messages obtained at sub-interface sid at the inside are output at the outside as being sent by the server in session sid via $\mathbf{SNET}_{N,\rho,n}$. Also, for clients $C \in \mathcal{C}$ marked as “active,” messages $m' \in \mathcal{M}$ obtained at the outside as input to $\mathbf{SNET}_{N,\rho,n}$ directed to C are input at sub-interface C at the inside.

We complete the proof by describing two reductions \mathbf{C} and \mathbf{C}' that connect with their inside interfaces to the games $\mathbf{G}^{\text{uf-cma}}$ and \mathbf{G}^{CR} , respectively. On a high level, the reductions emulate at the outside interface an emulation of $\mathbf{R} \doteq \prod_{C \in \mathcal{C}} \text{vrf}^C \text{sgn}^S[\mathbf{SNET}_{N,\rho,n}, \text{PKI}_{\mathfrak{F}}]$, where \mathbf{C} uses the key pair (sk, vk) obtained from $\mathbf{G}^{\text{uf-cma}}$. Thereby, \mathbf{C} computes the certificate $cert = (vk, f(vk))$ using the verification key vk obtained from $\mathbf{G}^{\text{uf-cma}}$, and the signatures $s \leftarrow \text{sign}(\eta_C | m'_{sid,1} | \eta_{sid} | m | cert; sk)$ in the converter sgn using signing queries to $\mathbf{G}^{\text{uf-cma}}$. A “forgery” in the emulated execution can then be used to win the game $\mathbf{G}^{\text{uf-cma}}$ (for the exact definition of forgery see the MBO below; the forgery for $\mathbf{G}^{\text{uf-cma}}$ is achieved by concatenating the corresponding nonces and the message). The reduction \mathbf{C}' emulates \mathbf{R} until a collision occurs in the hash function (this is easy to check because \mathbf{C}' knows all transmitted messages) and, in case it finds a collision, uses it to win the game \mathbf{G}^{CR} .

We define the MBOs \mathcal{E} and \mathcal{E}' on the systems \mathbf{R} , $\mathbf{S} \doteq \sigma^{E \rhd} \bullet'_{N,\rho,\mathfrak{F},SIG,n,h}$, and $\mathbf{C}'\mathbf{G}^{\text{CR}}$, which capture collisions in the hash functions as well as forgeries for the signature scheme. More concretely, \mathcal{E} becomes 1 once there are two different “truncated handshakes,” i.e. messages and nonces all visible at the E -interface, such that the evaluation of the hash function h on the two induced different inputs results in the same output. The MBO \mathcal{E}' is defined to become 1 once there is an input (\tilde{m}, \tilde{s}) at the E/sid -interface with $sid = (\eta_C, e)$ such that $\text{vrf}(\eta_C | \eta_{sid} | \tilde{m}, \tilde{s}; vk) = 1$ —such that η_{sid} was generated as a response in the S -session sid —unless there was an output (\tilde{m}, \tilde{s}') at some interface E/sid' with $sid' = (\eta, e')$ and $\eta_{sid'} = \eta_{sid}$ before. Then, proving “game equivalence” of the systems with the MBOs $\mathcal{E} \vee \mathcal{E}'$ and using Lemma 26 allows to conclude that the condition is fulfilled.

The equivalence can be seen as follows (we use the counter variables e_η in the same way they are defined in the systems, i.e., counting the number of sessions that have been initiated with nonce η):

- Initially, both systems output the certificate $cert = (vk, f(vk))$, with vk and f chosen according to the same distributions, as auxiliary information at the S -interface, and as an output corresponding to $\text{PKI}_{\mathfrak{F}}$ at the E -interface. Additionally, for each $C \in \mathcal{C}$, both systems output the client’s nonce η_C at the C -interface and forward messages from the C -interface to the E/C -sub-interface.
- Upon input (ack, η) at the outside E -interface corresponding to $\mathbf{SNET}_{N,\rho,n}$ and the delivery of the first message m' in session sid , the system \mathbf{R} outputs a nonce η_{sid} at the E -interface of $\mathbf{SNET}_{N,\rho,n}$ and the S/sid for $sid = (\eta, e_\eta)$. In \mathbf{S} , the nonce is generated according to the same distribution and output at the same interfaces. Subsequently, messages input at the E/sid -sub-interface are output at the S/sid -sub-interface.
- Upon the first message m input at the outside S/sid -interface with $sid = (\eta, e)$, system \mathbf{R} generates a signature s for $h(\eta_C | m' | \eta_{sid} | m | cert)$ using the key sk (within sgn), outputs s at the same sub-interface S/sid and the triple $(m, cert, s)$ at the E -interface of $\mathbf{SNET}_{N,\rho,n}$. Within \mathbf{S} , the signature is computed analogously by $\rhd \bullet_{N,\rho,\mathfrak{F},SIG,n}$, also output at S/sid , and the message/certificate/signature triple is output at the E/sid -sub-interface (with the same distribution) via σ .
- Upon input $(\text{deliver}, C, \tilde{\eta})$ at the outside E -interface corresponding to $\mathbf{SNET}_{N,\rho,n}$, there is no immediate output (neither for \mathbf{R} nor for \mathbf{S}).
- Upon delivering the first message \tilde{m}_1 via $\mathbf{SNET}_{N,\rho,n}$ to a client C , if $\tilde{m}_1 \neq (m, cert, s)$ output at the E -sub-interface corresponding to the same session before, or if the server’s

nonce was not delivered faithfully in the same session (note that the client nonce defines the session), then, in **R**, **vrf** aborts since: either \tilde{m}_1 is not a well-formed certificate or the verification at $\text{PKI}_{\mathfrak{F}}$ fails, or the verification of the signature fails—this is guaranteed by the fact that otherwise either a collision in the hash function or a forgery of the signature would occur. In **S**, σ marks C as “failed” in the same cases. There is no output, neither in **R** nor in **S**.

- After a certain session (either $C \in \mathcal{C}$ or sid at S) has been initialized, messages input there are output at the corresponding sub-interface of the E -interface and vice versa. This is consistent in both **R** and **S**.

As the same arguments hold for the case $\mathbf{C}'\mathbf{G}^{\text{CR}}$, and with the exception that the signatures are obtained from $\mathbf{G}^{\text{uf-cma}}$ but have the same distribution also in the case $\mathbf{C}\mathbf{G}^{\text{uf-cma}}$, and each violation of $\mathcal{E} \vee \mathcal{E}'$ can be used to win either \mathbf{G}^{CR} or $\mathbf{G}^{\text{uf-cma}}$, this concludes the proof. \square

Constructing the premaster secret. The construction step (we first only construct the premaster secret) is then achieved by the protocol (**dhe13c**, **dhe13s**), in which the server chooses for each session a (potentially fresh) Diffie-Hellman group and element, which are sent as an authenticated message via $\succsim \bullet'$. We denote the group used by the protocol by \mathbb{G} .¹⁷

The distribution AUX in this case consists of two parts. The first part is the same for all sessions and consists of the certificate (i.e., depends on \mathfrak{F} and SIG of $\succsim \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n, h}$). The distribution is described by $(sk, vk) \leftarrow \text{gen}$, $f \leftarrow \mathfrak{F}$, and then $\text{cert} = (vk, f(vk))$. The second part is chosen independently for each session by doing the same process as in the protocol: choose two group elements $g_1, g_2 \leftarrow \mathbb{G}$ uniformly at random.

The client’s converter **dhe13c** initially obtains η_C at the inside interface, and then:

- Choose $u \leftarrow \{1, \dots, q\}$ (with $q = |\mathbb{G}|$) and input g^u at the inside interface.
- Receiving (aux, η_{sid}, m, s) , parse the message as $g' = m$ (abort if impossible; note that $m \in \{0, 1\}^*$ while $g \in \mathbb{G}$). Output $(g'^u, \eta_C, \eta_{sid}, aux|g^u|m|s)$.
- Forward the following communication between the inside and the outside interfaces.

The server’s converter **dhe13s** connects to the S -interface of $\succsim \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n, h}$. Both the inside and outside interfaces have sub-interfaces $\text{sid} = (\eta_C, e) \in \mathcal{N} \times [n]$. The converter behaves as follows:

- Initially, receive aux on the inside interface.
- Upon obtaining a nonce η_{sid} and a group element \tilde{g} at sub-interface $\text{sid} = (\eta_C, e)$, choose an exponent $v \leftarrow \{1, \dots, |\mathbb{G}|\}$. Send $m = g^v$ via the inside sub-interface sid . Obtain the signature s at the inside interface in return. Output $(\tilde{g}^v, \eta_{sid}, aux|\tilde{g}|m|s)$ at the outside sub-interface sid .
- Forward the following communication between the corresponding sub-interfaces sid of the inside and the outside interface.

The described protocol indeed constructs the master secret key resource from the network $\succsim \bullet'$ under the DDH assumption in the group \mathbb{G} .

Lemma 14. *The protocol (**dhe13c**, **dhe13s**) constructs from the assumed resource $\succsim \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n, h}$ the resource $\text{KEY}_{\mathbf{N}, \rho, AUX, n, \mathcal{K}}$ with $\mathcal{K} = \mathbb{G}$, under the DDH assumption for $\mathbb{G} = \langle g \rangle$. More formally, for the simulator σ and the reduction \mathbf{C}_1 and \mathbf{C}_2 described in the proof,*

$$\succsim \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, SIG, n, h} \xrightarrow{(\text{dhe13c}, \text{dhe13s}), \sigma, (\varepsilon_1, \varepsilon_2)} \text{KEY}_{\mathbf{N}, \rho, AUX, n, \mathcal{K}},$$

¹⁷For now, we ignore the effects of different groups being available.

with $\varepsilon_i(\mathbf{D}) \doteq n \cdot |\mathcal{C}| \cdot \Delta^{\mathbf{DC}_i}((g^A, g^B, g^{AB}), (g^A, g^B, g^C))$, with A, B, C uniformly random from the set $\{0, \dots, |\mathbb{G}| - 1\}$, for all distinguishers \mathbf{D} .

Proof.

Availability. We first show the availability condition. The “ideal” system $\perp^E \text{KEY}_{\mathbf{N}, \rho, \text{AUX}, n, \mathcal{K}}$ chooses, for each session $sid = (\eta_C, 1)$ with $\rho(C) = \eta_C$, a nonce η_{sid} , a uniformly random key $\kappa_C \in \mathcal{K}$ and auxiliary information aux_C (which includes uniformly random group elements from the key exchange), and outputs $(\kappa_C, \eta_{sid}, aux_C)$, and afterward forwards communication between C and S/sid .

The distribution in the case $\prod_{C \in \mathcal{C}} \text{dhe13c}^C \text{dhe13s}^S \perp^E \left(\succ \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n, h} \right)$ is as follows. The key κ_C for each client C is a group element in \mathbb{G} which completes the Diffie-Hellman triples with the group elements in the auxiliary information, the nonces η_{sid} for $sid = (\eta_C, 1)$ are distributed according to \mathbf{N} , and the auxiliary information aux_C has exactly the same distribution as well. Moreover, after at both interfaces C and S/sid the above information is output, the resource behaves as a bidirectional channel between those interfaces.

The proof is basically completed via a hybrid argument, where the sessions are ordered in some particular way (e.g., by taking the lexicographic ordering on \mathcal{C} and the natural one on the sessions per associated nonce). The reduction \mathbf{C}_1 chooses one particular session for embedding the challenge (the first two group elements as those exchanged in the protocol, all “smaller” sessions are served with freshly chosen Diffie-Hellman triples, and all “greater” sessions are served by purely random triples. A standard hybrid argument completes the proof of the availability statement.

Security. The simulator basically needs to take care of the $\succ \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n, h}$ ’s E -interface (beginning with choosing some good aux), and then needs to simulate the DH exchange, i.e., the group elements generated by the client and the server. The simulator σ behaves as follows:

- Throughout, σ keeps counters e_η for each $\eta \in \mathcal{N}$ in the usual way (i.e., increase e_η whenever the nonce η is delivered to the server).
- Initially, obtain a string $cert \in \{0, 1\}^*$ at the outside interface (this is needed for the auxiliary information).
- For each $C \in \mathcal{C}$, sample a group element $\bar{g}_{sid} = g^u$ for a uniformly random $u \in \{1, \dots, |\mathbb{G}|\}$ and output \bar{g}_{sid} at the E/C -sub-interface.
- Upon input (ack, η, m') at the outside, where m' is a valid group element (otherwise abort), input (ack, η) at the inside and obtain as response η_{sid} for $sid = (\eta, e_\eta)$. Output η_{sid} at the outside interface. Also, choose a value $v \leftarrow_{\$} \{1, \dots, |\mathbb{G}|\}$, compute $\tilde{g}_{sid} = g^v$ and $s \leftarrow \text{sign}(h(\eta|m'|_{\eta_{sid}}|\tilde{g}_{sid}|cert); sk)$, and output (\tilde{g}_{sid}, s) at the outside sub-interface sid . Furthermore
 - if $m' = \bar{g}_{sid}$, then input $(\kappa_S, \eta_C, e, aux)$ and $aux = \bar{g}_{sid}|cert|\tilde{g}_{sid}|s$ (with $sid = (\eta_C, e)$) at the inside interface;
 - otherwise, input $(\text{inject}, \eta_C, e, aux, g'_{sid}{}^v)$ at the inside interface with $sid = (\eta_C, e)$ and $aux = m'|cert|\tilde{g}_{sid}|s$.
- Upon input $(\text{deliver}, C)$ at the outside, if $(\text{ack}, \eta_C, \bar{g}_{sid})$ was input before, and there is a session $sid = (\eta_C, e)$ with $\tilde{\eta} = \eta_{sid}$, then input $(\text{key-c}, C, aux, \eta_{sid})$ with $aux = \bar{g}_{sid}|cert|\tilde{g}_{sid}|s$.
- Deliver messages faithfully for the sessions where the setup is complete.

We first observe that the two systems $\mathbf{R} = \prod_{C \in \mathcal{C}} \text{dhe13c}^C \text{dhe13s}^S \left(\succsim \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n, h} \right)$ and $\mathbf{S} = \sigma^E \text{KEY}_{\mathbf{N}, \rho, \text{AUX}, n, \mathcal{K}}$ are equivalent with respect to how they treat nonces; in particular, the responses to $(\text{ack}, \eta, *)$ are determined by choosing fresh parameters in both cases. The same argument holds for the queries in which messages are forwarded in the sessions that completed the setup. For the group elements in a particular session sid , the difference between \mathbf{R} and \mathbf{S} is that in \mathbf{R} , all “keys” output at either $C \in \mathcal{C}$ or $S/(\eta_C, e)$ for $\eta_C \in \mathcal{N}$ and $e \in [n]$ are consistently computed Diffie-Hellman keys, whereas in \mathbf{S} they are uniformly random and independent of the group elements simulated at the E -interface and included in the auxiliary information.

We then describe a reduction system \mathbf{C}_2 , which obtains at the inside interface three group elements $\bar{g}, \tilde{g}, \hat{g} \in \mathbb{G}$. We assume that there is some (e.g., lexicographic) ordering on the set \mathcal{C} , i.e., the elements are $C_1, \dots, C_{|\mathcal{C}|}$. The system \mathbf{C}_2 then behaves as follows (we stress that every query can be associated with a (client’s) nonce $\eta \in \mathcal{N}$, either because the nonce is given explicitly, or because the query belongs to a session that is described by a nonce and a counter):

- In the beginning, \mathbf{C}_2 chooses uniformly at random a session to embed the challenge (i.e., it chooses a client $\bar{C} \in \mathcal{C}$ and a session number \bar{e} at the server). Then, for \bar{C} , output \bar{g} at the E/\bar{C} -sub-interface, and for each $C \in \mathcal{C} \setminus \{\bar{C}\}$, it chooses a value $r_C \in \{0, \dots, |\mathbb{G}| - 1\}$ and outputs g^{r_C} at the E/C -sub-interface.
- For queries (at the S - and E -interfaces) that are related to nonces $\eta \in \mathcal{N} \setminus \rho(\mathcal{C})$, the system \mathbf{C}_2 can easily reproduce the behavior of the real or ideal systems (their behavior is equivalent for those queries).
- For a $(\text{ack}, \eta_C, \bar{g}_C)$ -query (i.e., $sid = (\eta_C, e)$ for some $e \in \mathbb{N}$, and $\bar{g}_C \in \mathbb{G}$ —otherwise return nothing):
 - If \bar{g}_C is different from the group element chosen in that session before, choose a random value $r_{sid} \in \{0, \dots, |\mathbb{G}| - 1\}$, compute the group element $\tilde{g}^{r_{sid}}$ and the signature $s_{sid} = \text{sign}(h(\eta_C | \bar{g}_C | \eta_{sid} | g^{r_{sid}} | \text{cert}); sk)$ and output $(g^{r_{sid}}, s_{sid})$ at the E/sid -sub-interface. Output $\tilde{g}^{r_{sid}}$ at the S/sid -sub-interface. Otherwise, if the group element is forwarded faithfully:
 - If the session is “smaller” than the one chosen for embedding the challenge, then (as above) a random value is chosen and the response and key are computed as above.
 - If $C = \bar{C}$ and $e = \bar{e}$, and if $\bar{g}_C = \bar{g}$, then use $\tilde{g}_C = \tilde{g}$ for the response and $\kappa_C = \hat{g}$ for the key.
 - If the session is “greater” than the one chosen for embedding the challenge, then use a purely random group element for the response, and another uniformly random group element for the key.
- For the query $(\text{deliver}, C)$ at the E -sub-interface corresponding to $\succsim \bullet'_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n, h}$: if the earlier transmission of the client’s group element together with the nonce has been unmodified, then output the same key as in the server’s session.

The proof is concluded via a standard hybrid argument. □

Constructing the handshake master secret. In the following step, we use the PMK generated above and derive from it first the handshake master secret. The protocol achieving this is as follows: Both client and server compute

$$hs_ms = \text{eval}_{\text{PRF}}(pms, \text{handshake master secret}, \text{session_hash}).$$

Here, pms is obtained from the previous $\text{KEY}_{N,\rho,AUX,n,\mathcal{K}}$ resource, and $session_hash$ can be computed from the auxiliary information. We denote the converters by hsc and hss , respectively.

Lemma 15. *The protocol (hsc, hss) constructs from the assumed resource $\text{KEY}_{N,\rho,AUX,n,\mathcal{K}}$ with $\mathcal{K} = \mathbb{G}$ the resource $\text{MSK}_{N,\rho,AUX,n}$, given that the hash function is collision resistant and HMAC is a PRF when keyed with string representing a uniformly random element from \mathbb{G} . More formally, for the simulator σ and the reduction \mathbf{C} and \mathbf{C}' described in the proof,*

$$\text{KEY}_{N,\rho,AUX,n,\mathcal{K}} \xrightarrow{(\text{hsc}, \text{hss}), \sigma, (\varepsilon, \varepsilon)} \text{MSK}_{N,\rho,AUX,n},$$

such that for all distinguishers \mathbf{D} , $\varepsilon(\mathbf{D}) \doteq \Gamma^{\mathbf{DC}}(\mathbf{G}^{\mathbf{CR}}) + |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'}(\text{prf } \mathbf{U}_{\mathbb{G}}, \mathbf{F})$.

Proof sketch. The availability and security cases are almost the same (the only interesting cases are where a key not known to the attacker is expanded; the other cases can be treated easily). If the session hash values used as input to the PRF are all distinct, and the keys to the PRF are uniformly chosen, then the outputs of the PRFs are distinct samples. This is formalized by a condition that corresponds to a collision in the output of a session hash (which can be computed from the auxiliary information; also a reduction is easy to design), and then by a hybrid argument replacing one-by-one for each $C \in \mathcal{C}$ the output of the PRF by a uniformly random value in $\{0, 1\}^{384}$. \square

4 Expanding the Key

The keys constructed in Section 3 are not directly used in the record layer protocols. In TLS 1.2, the constructed master secret is expanded (using a PRF) to obtain keys for encryption and MAC of the record protocol, as well as for the computation of the “finished” messages. In TLS 1.3, the constructed handshake master secret is expanded (again, using a PRF) to obtain the finished messages, keys for the authenticated encryption used to encrypt those messages, and the actual master secret. This secret is then further expanded to obtain keys for the authenticated encryption used to protect the actual payload messages.

The encryption of the “finished” messages is performed differently in TLS 1.2 and TLS 1.3. In particular, TLS 1.2 uses the record layer protocol with the same keys used to protect the payload messages. In contrast, TLS 1.3 uses specific handshake keys derived from the handshake master secret. To account for these differences, we have the resources constructed in this section output the “finished” messages in the form where they are already protected by the record layer scheme. This allows us to unify the subsequent protocol steps.

4.1 Key Expansion in TLS 1.2

The key expansion step in TLS 1.2 is described as a protocol $(\text{expc}, \text{exps})$, which uses a HMAC-based PRF to generate sufficient key material (depending on the actual cipher suite) for two encryption and two MAC keys (one key per purpose and direction). Furthermore, the converters also generate the so-called “finished” messages which are used by the client and the server to confirm the computed keys.

More formally, starting from the resource $\text{MSK}_{N,\rho,AUX,n}$, we expand the key using a pseudo-random function (PRF); our goal is to construct the ideal resource $\bigotimes_{C \in \mathcal{C}} \llbracket \overset{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(C, S/\rho(C))}$. An advantage of this description is that it is the parallel composition of multiple “single-client” resources $\overset{KSP,*}{=} \bullet_{\text{cphs},n}$, which means all following protocol steps can be proven in a setting where there is only a single client, and then composed using the generic composition theorem.

The (extended) key space actually depends on the cipher suites in use. In this work, we do not specifically focus on any of the cipher suites, but rather define a function $\text{cphs} : \{\text{CAuth}, \text{SAuth}, \text{CEnc}, \text{SEnc}, \text{CIV}, \text{SIV}\} \rightarrow \mathbb{Z}$, which outputs the length of each of the following keys: $\text{client_write_MAC_key}$ ($\kappa_{C,a}$), $\text{server_write_MAC_key}$ ($\kappa_{S,a}$), client_write_key ($\kappa_{C,e}$), server_write_key ($\kappa_{S,e}$), client_write_IV ($\kappa_{C,IV}$), server_write_IV ($\kappa_{S,IV}$). Note that the last two keys are often not generated, since they are only used for implicit nonce techniques, see [DR08]. Usually, the key length for each of the first four keys is 32 bytes, i.e. 256 bits. By convention, if $\kappa_{C,IV}, \kappa_{S,IV}$ are not generated, we write that $\text{cphs}(\kappa_{C,IV}) = \text{cphs}(\kappa_{S,IV}) = 0$. The set of all possible keys (parsed as a concatenation of all the aforementioned keys, in the given order) is denoted \mathcal{K} . The derived keys for each sessions are, in this sequence, $\text{client_write_MAC_key}$, $\text{server_write_MAC_key}$, client_write_key , server_write_key , client_write_IV , server_write_IV , and the two “finished” messages which are of length 96 bits each.

We usually write the generated vectors as $\vec{\kappa} = (\kappa_{C,a}, \kappa_{S,a}, \kappa_{C,e}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV}, \xi_C, \xi_S)$, with the lengths of the individual components indicated by cphs .

$$\begin{aligned} &KSP,* \\ &= \bullet_{\text{cphs},n} \end{aligned}$$

Set $b_e = 0$ for $e \in [n]$ and $b_C = 0$.

No attacker present: Output at both interfaces C and $S/1$ the same vector $\vec{\kappa}$ described above. Subsequently, provide bidirectional channels between those two interfaces.

Attacker present:

- Upon input of the type $(\text{key-c}, e)$ with $e \in [n] \cup \{0\}$ at the E -interface, if $b_C = 0$, then:^a
 1. If $\vec{\kappa}^e$ is undefined, draw $\vec{\kappa}^e$ at random as described above.
 2. Output $\vec{\kappa}^e$ at the C -interface and set $b_C = 1$.

Then, relay all communication between the interfaces C and E/C .

- On input (deliver) at the E/e -sub-interface, if $b_e = 0$, then:
 1. If $\vec{\kappa}^e$ is undefined, draw $\vec{\kappa}^e$ at random as described above.
 2. Output $\vec{\kappa}^e$ at the S/e -interface and set $b_e = 1$.

Then, relay all communication between the interfaces S/e and E/e .

- Upon input $(\text{inject}, e, \vec{\kappa})$ with $e \in [n]$, if $b_e = 0$ then output $\vec{\kappa}$ at interface S/e and set $b_e = 1$. Then, relay all communication between the interfaces S/e and E/e .

^aThe value $e = 0$ allows for client sessions that are not paired with any server session.

In TLS the session keys are obtained via a PRF based on HMAC, taking as input the master secret value obtained in a previous step. Afterward, the client and server each generates a final message by again querying the PRF keyed with the master secret value, on input the hash of a concatenation of messages (basically the transcript of the session). This hash function, denoted H , is required to be collision resistant. For the purpose of our analysis, we will assume that the key material and the finished messages are derived by means of a pseudo-random function $\text{PRF} = (\text{gen}_{\text{PRF}}, \text{eval}_{\text{PRF}})$, with output length equal to $\max(|\text{cphs}(\cdot)|, 96)$. For a more detailed discussion how this is achieved in TLS, we refer the reader to Appendix A.3.

The client converter, expc , behaves as follows:

1. Obtain the values $(\kappa, \eta_C, \eta_{sid}, m)$ from $\text{MSK}_{N,\rho,AUX,n}$, where m is a concatenation of messages.
2. Generate keys $(\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV}) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion}|\eta_C|\eta_{sid})$.
3. Generate messages $(\xi_C, \delta_C) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished}|H(\eta_C|\eta_{sid}|m|\gamma))$ and $(\xi_S, \delta_S) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished}|H(\eta_C|\eta_{sid}|m|c_{\xi_C}|\gamma))$.¹⁸ In the above computation, the constant γ stands for the “ChangeCipherSpec” message, whereas the value c_{ξ_C} is computed as a function of ξ_C in a way that depends, as do the lengths of the computed keys, on the adopted cipher suite (this corresponds to the encryption of the finished message, computed in the record layer protocol).
4. Compute $\bar{\xi}_C$ and $\bar{\xi}_S$ by protecting ξ_C and ξ_S using the keys $\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}$ and encryption and authentication as specified by the cipher suite.
5. Output $(\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV})$ and $(\bar{\xi}_C, \bar{\xi}_S)$.

The server converter, **exps** behaves as follows. For each of the sessions described by a pair $sid = (\eta_C, e) \in \mathcal{N} \times [n]$:

1. Obtain the values (κ, η_{sid}, m) from $\text{MSK}_{N,\rho,AUX,n}$, where m is a concatenation of messages.
2. Generate keys $\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV} \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion}|\eta_C|\eta_{sid})$.
3. Generate “finished” messages $(\xi_C, \delta_C) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished}|H(\eta_C|\eta_{sid}|m|\gamma))$ and $(\xi_S, \delta_S) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished}|H(\eta_C|\eta_{sid}|m|c_{\xi_C}|\gamma))$.
4. Compute $\bar{\xi}_C$ and $\bar{\xi}_S$ by protecting ξ_C and ξ_S using the keys $\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}$ and encryption and authentication as specified by the cipher suite.
5. Output $(\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV})$ and $(\bar{\xi}_C, \bar{\xi}_S)$.

We aim for the statement that the protocol $(\text{expc}, \text{exps})$ constructs from $\text{MSK}_{N,\rho,AUX,n}$ the parallel composition of one copy of $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$ for each client $C \in \mathcal{C}$, each such resource corresponding to a server’s sub-interface S/η for the nonce $\eta = \rho(C)$ chosen by client C . This, however, is not directly true since by the definition of the server’s converter **exps**, an attacker can generate output at a server’s sub-interface S/η for client nonces $\eta \notin \rho(\mathcal{C})$. This is not an artifice of our constructive treatment, but an inherent trait of protocols with unilateral authentication: the key may be shared either between clients and servers *or* between attackers and servers. We capture this feature at the level of the key expansion step by introducing “residual” resources $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$ for which the client’s interface is also provided as a sub-interface of the E -interface.

We denote the parallel composition of these resources as $\bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)}$, where the double-bracket notation indicates that the interfaces of each one of the (three-interface) resource $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$ are provided as interfaces I and J for each $(I, J) \in \mathcal{P}$.

We show the following result, where we describe the pseudo-random function PRF by a converter **prf** and denote the system outputting a single 384-bit random string by \mathbf{U}_{384} and the uniform random function (with the same output length as PRF) by \mathbf{F} .

Lemma 16. *The protocol $(\text{expc}, \text{exps})$ constructs $\bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)}$ from $\text{MSK}_{N,\rho,AUX,n}$, for $\mathcal{P} = \{(C, S/\rho(C)) : C \in \mathcal{C}\} \cup \{(E/\eta, S/\eta) : \eta \notin \rho(\mathcal{C})\}$. The construction holds under the assumption that the hash function H is collision resistant and the pseudo-random function PRF is indistinguishable from a random function. In more detail, for the simulator σ and the reductions \mathbf{C} and \mathbf{C}' described in the proof,*

$$\text{MSK}_{N,\rho,AUX,n} \xRightarrow{(\text{expc}, \text{exps}), \sigma, (\varepsilon, \varepsilon)} \bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)}$$

¹⁸The “extra” bits δ_C and δ_S are discarded.

such that for all distinguishers \mathbf{D} , $\varepsilon(\mathbf{D}) \preceq \Gamma^{\mathbf{DC}}(\mathbf{G}^{\mathbf{CR}}) + |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$.

Proof.

Availability. We show the availability condition. Let $\mathbf{R}_\perp \doteq \prod_{C \in \mathcal{C}} \text{expc}^C \text{exp}^S \perp^E \text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ and $\mathbf{S}_\perp \doteq \perp^E \left[\bigotimes_{(I, J) \in \mathcal{P}} \llbracket \stackrel{KSP, *}{=} \bullet_{\text{cphs}, n} \rrbracket^{(I, J)} \right]$. We notice that the main difference between the real and ideal system (when the cheating bit is set to $b = 0$), is that the latter chooses the keys and the finished messages (ξ_C, ξ_S) to be output at C and $S/1$ uniformly at random, whereas the former computes such values via the PRF.

We introduce a first hybrid systems \mathbf{H}_\perp , to deal with possible collisions in the hash function H . The system \mathbf{H}_\perp is defined exactly as \mathbf{R}_\perp , with one difference: Whenever there exists two distinct tuples $(\eta_C, \eta_{sid}, m, c_{\xi_C})$ and $(\eta'_C, \eta'_{sid}, m', c'_{\xi'_C})$ such that either $(\eta_C, \eta_{sid}, \gamma)$ and $(\eta'_C, \eta'_{sid}, \gamma)$ or $(\eta_C, \eta_{sid}, m, c_{\xi_C}, \gamma)$ and $(\eta'_C, \eta'_{sid}, m', c'_{\xi'_C}, \gamma)$ are a collision for H , then the corresponding output of the hash function is re-sampled uniformly until a completely fresh (i.e., not previously used) value is found. We argue that a distinguisher between the two systems \mathbf{R} and \mathbf{H}_\perp can be used to build a reduction \mathbf{C} breaking collision resistance of H . The reduction \mathbf{C} connects with the inside interface to the game $\mathbf{G}^{\mathbf{CR}}$ and provides at the outside interface an emulation of \mathbf{R}_\perp , using a description of the hash function obtained from $\mathbf{G}^{\mathbf{CR}}$. A “collision” in the emulated execution can then be used to win the game $\mathbf{G}^{\mathbf{CR}}$ (for the exact definition of collision see the MBO below).

We define the monotone binary output (MBO) \mathcal{E} on the systems \mathbf{R} , \mathbf{H}_\perp , and $\mathbf{CG}^{\mathbf{CR}}$ as the following “collision” event: it becomes 1 once there are two distinct tuples $(\eta_C, \eta_{sid}, m, c_{\xi_C})$ and $(\eta'_C, \eta'_{sid}, m', c'_{\xi'_C})$ such that either $(\eta_C, \eta_{sid}, \gamma)$ and $(\eta'_C, \eta'_{sid}, \gamma)$ or $(\eta_C, \eta_{sid}, m, c_{\xi_C}, \gamma)$ and $(\eta'_C, \eta'_{sid}, m', c'_{\xi'_C}, \gamma)$ are a collision for H . Clearly the random systems \mathbf{R} and \mathbf{H}_\perp are equivalent conditioned on the MBO not being 1. Thus invoking Lemma 26 allows to conclude that $\Delta^{\mathbf{D}}(\mathbf{R}, \mathbf{H}_\perp) \leq \Gamma^{\mathbf{DC}}(\mathbf{G}^{\mathbf{CR}})$.

Next, we argue that a distinguisher between \mathbf{H}_\perp and \mathbf{S}_\perp can be used together with a reduction \mathbf{C}' to distinguish the pseudo-random function PRF from a truly random function. The reduction is actually a series of reductions \mathbf{C}'_i for $i \in [|\mathcal{C}|]$, where the index pinpoints one client $C \in \mathcal{C}$. We assume some (e.g. lexicographic) order \preceq over the set of clients, and write C_i to denote the i -th client with respect to this order. The reduction \mathbf{C}'_i makes 3 queries to connected system and works as follows:

- For all the sessions between a client $C \succ C_i$, behave as in \mathbf{H}_\perp .
- For all the sessions between a client $C \prec C_i$, behave as in \mathbf{S}_\perp .
- In the session between C_i and $sid = (\eta, 1)$, forward $x_1 = \text{key expansion}|\eta|\eta_{sid}$, $x_2 = \text{client finished}|H(\eta|\eta_{sid}|m|\gamma)$ and $x_3 = \text{server finished}|H(\eta|\eta_{sid}|m|c_{x_2}|\gamma)$ to the connected system; note that the value c_{x_2} can be computed as a function of x_2 by only knowing the cipher suite in use. Emulate the session using the corresponding values y_1 , y_2 and y_3 , received from the inside interface. (In case a collision is found, similar to the MBO defined in \mathbf{H}_\perp , re-sample the output of the hash function uniformly.)

We note that if \mathbf{C}'_i is connected to \mathbf{F} , then the values y_1 , y_2 and y_3 are uniform, whereas if \mathbf{C}'_i is connected to $\text{prf } \mathbf{U}_{384}$, then they are computed as $\text{eval}_{\text{PRF}}(\kappa, x)$ for $x \in \{x_1, x_2, x_3\}$ and a uniformly random key κ . Also, for all $0 \leq i \leq |\mathcal{C}|$, $\mathbf{C}'_i \text{prf } \mathbf{U}_{384} \equiv \mathbf{C}'_{i+1} \mathbf{F}$. Furthermore, as $\mathbf{C}'_0 = \mathbf{H}_\perp \text{prf } \mathbf{U}_{384}$ and $\mathbf{C}'_{|\mathcal{C}|} \mathbf{F} = \mathbf{S}_\perp$, the above argument concludes the proof of the availability condition, where \mathbf{C}' chooses any one of the \mathbf{C}'_i uniformly at random.

Security. To prove the security condition, consider the following simulator σ :

- Initially, σ sets $e_\eta = 1$ for all $\eta \in \mathcal{N}$. (The counters are kept consistent, i.e. they are increased whenever the simulator receives at the outside interface an **(ack, *)** or **(inject, *)** command.)
- Upon input **(ack, η)** at the outside interface, if $e_\eta \leq n$, then choose η_{sid} for $sid = (\eta, e_\eta)$ and output η_{sid} at the outside interface as being transmitted via $\text{MSK}_{N,\rho,AUX,n}$.
- Upon input **(key-c, $C, aux, \tilde{\eta}$)** at the outside interface (the first time for this $C \in \mathcal{C}$):
 - If $\tilde{\eta} = \eta_{sid}$ for $sid = (\rho(C), e')$ for some $e' < e_{\rho(C)}$, let $e = e'$,
 - else, let $e = 0$.

Issue **(key-c, e)** at the C -sub-interface of $\bigotimes_{C \in \mathcal{C}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(C,S/\rho(C))}$. Afterward, forward communication between the outside and the inside C -sub-interfaces.

- Upon input **(deliver, η, e, aux)**, if $\eta \in \rho(\mathcal{C})$, $e < e_\eta$, issue **(deliver)** at the $E/\rho^{-1}(\eta)$ -sub-interface of the system $\bigotimes_{C \in \mathcal{C}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(C,S/\rho(C))}$.
- Upon input **(inject, η, e, aux, κ)** at the outside interface, define $sid = (\eta, e)$ and compute

$$\begin{aligned} (\tilde{\kappa}_{C,a}, \tilde{\kappa}_{S,a}, \tilde{\kappa}_{C,e}, \tilde{\kappa}_{S,e}, \tilde{\kappa}_{C,IV}, \tilde{\kappa}_{S,IV}) &\leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion}|\eta|_{\eta_{sid}}) \\ \tilde{\xi}_C &\leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished}|H(\eta|_{\eta_{sid}}|aux|\gamma)) \\ \tilde{\xi}_S &\leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished}|H(\eta|_{\eta_{sid}}|aux|c_{\tilde{\xi}_C}|\gamma)). \end{aligned}$$

Then issue **(inject, $e, \tilde{\kappa}_{C,a}, \tilde{\kappa}_{S,a}, \tilde{\kappa}_{C,e}, \tilde{\kappa}_{S,e}, \tilde{\kappa}_{C,IV}, \tilde{\kappa}_{S,IV}, \tilde{\xi}_C, \tilde{\xi}_S$)** at the inside $\rho^{-1}(\eta)$ -sub-interface of $\bigotimes_{C \in \mathcal{C}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(C,S/\rho(C))}$.

- After either a **(deliver, $\eta, e, *$)** or an **(inject, $\eta, e, *, *$)**, forward communication between the outside and inside $sid = (\rho(\eta), e)$ sub-interfaces.

For the sake of brevity, we use the notation $\mathbf{R} \doteq \prod_{C \in \mathcal{C}} \text{exp}_C^C \text{exp}_S^S \text{MSK}_{N,\rho,AUX,n}$ as well as $\mathbf{S} \doteq \sigma^E \left[\bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)} \right]$. We notice that the only difference between \mathbf{R} and \mathbf{S} is within **(key-c, *)** commands, as in the former the key material and the pair of values (ξ_C, ξ_S) are computed via the PRF and the hash function, whereas in the latter they are sampled uniformly. Similar to the availability proof, we consider a hybrid system \mathbf{H} where we re-sample outputs of H corresponding to distinct inputs generating a collision (until a “fresh” value is found). Then one can describe a reduction \mathbf{C}'' that, together with a distinguisher telling apart \mathbf{R} and \mathbf{H} , breaks the collision resistance of H . The description of the reduction is similar to the one considered in the availability case, and is therefore omitted.

Finally, we argue that a distinguisher between \mathbf{H} and \mathbf{S} can be used to build a reduction \mathbf{C}''' breaking the pseudo-randomness of the PRF. We remark that there is a single key κ_C associated with client C , and the key at some interface $S/(\rho(C), e)$ is either the same key κ_C or it is an injected key. It follows that the the description of the reduction goes along the same lines to the one considered for the availability proof. Put together, the above arguments conclude the proof of the security condition. \square

4.2 Key Expansion in TLS 1.3

The expansion in TLS 1.3 is computed differently. The underlying resource was defined to output the handshake master secret, and hence the computations performed here are as follows:

- Compute the handshake session keys and the handshake finished message (encrypted with the session keys—note that the resource outputs the “encrypted” version of the messages),
- compute the “main” master secret,
- compute the session keys for the main protocol,
- output the encrypted finished messages and all the session keys.

The next protocol step, which we describe as a protocol (`expc13`, `exps13`), uses a HMAC-based PRF to generate sufficient key material (depending on the actual cipher suite) for two encryption and two MAC keys (one key per purpose and direction). Furthermore, the converters also generate the so-called “finished” messages which are used by the client and the server to confirm the computed keys.

More formally, starting from the resource $\text{MSK}_{N,\rho,AUX,n}$, we expand the key using a pseudo-random function (PRF); our goal is to construct the ideal resource $\bigotimes_{C \in \mathcal{C}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(C,S/\rho(C))}$. An advantage of this description is that it is the parallel composition of multiple “single-client” resources $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$, which means all following protocol steps can be proven in a setting where there is only a single client, and then composed using the generic composition theorem.

In TLS 1.3 the session keys are obtained via a PRF based on HMAC, taking as input the handshake master secret value obtained in a previous step. Additionally, the client and server each generates a finished message by again querying the PRF keyed with the handshake master secret value, on input the hash of a concatenation of messages (basically the transcript of the session). This hash function, denoted H , is required to be collision resistant. For the purpose of our analysis, we will assume that the key material and the finished messages are derived by means of a pseudo-random function $\text{PRF} = (\text{gen}_{\text{PRF}}, \text{eval}_{\text{PRF}})$, with output length equal to $\max(|\text{cphs}(\cdot)|, 96)$. For a more detailed discussion how this is achieved in TLS, we refer the reader to Appendix A.3.

The client converter, `expc13`, behaves as follows:

1. Obtain the values $(\kappa, \eta_C, \eta_{sid}, m)$ from $\text{MSK}_{N,\rho,AUX,n}$, where m is a concatenation of messages.
2. Compute the session hash $h = H(m)$.
3. Compute handshake session keys for the authenticated encryption (from the handshake master secret) as $(\hat{\kappa}_C, \hat{\kappa}_S, \hat{\kappa}_{C,IV}, \hat{\kappa}_{S,IV}) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion}|\eta_C|\eta_{sid})$.
4. Generate a master secret key $\bar{\kappa} \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{extended master secret}|h)$.
5. Generate keys $(\kappa_C, \kappa_S, \kappa_{C,IV}, \kappa_{S,IV}) \leftarrow \text{eval}_{\text{PRF}}(\bar{\kappa}, \text{key expansion}|\eta_C|\eta_{sid})$.
6. Generate confirmation messages $(\xi_C, \delta_C) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished}|h)$ and $(\xi_S, \delta_S) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished}|h)$.¹⁹
7. Compute $\bar{\xi}_C$ and $\bar{\xi}_S$ by protecting ξ_C and ξ_S using the keys $\hat{\kappa}_C, \hat{\kappa}_S, \hat{\kappa}_{C,IV}, \hat{\kappa}_{S,IV}$ and authenticated encryption as specified by the cipher suite.
8. Output $(\kappa_C, \kappa_S, \kappa_{C,IV}, \kappa_{S,IV})$ and $(\bar{\xi}_C, \bar{\xi}_S)$.

The server converter, `exps13` behaves as follows. For each of the sessions described by a pair $sid = (\eta_C, e) \in \mathcal{N} \times [n]$:

¹⁹The “extra” bits δ_C and δ_S are discarded.

1. Obtain the values (κ, η_{sid}, m) from $\text{MSK}_{N,\rho,AUX,n}$, where m is a concatenation of messages.
2. Compute the session hash $h = H(m)$.
3. Compute handshake session keys for the authenticated encryption (from the handshake master secret) as $(\hat{\kappa}_C, \hat{\kappa}_S, \hat{\kappa}_{C,IV}, \hat{\kappa}_{S,IV}) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion}|\eta_C|\eta_{sid})$.
4. Generate a master secret key $\bar{\kappa} \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{extended master secret}|h)$.
5. Generate keys $(\kappa_C, \kappa_S, \kappa_{C,IV}, \kappa_{S,IV}) \leftarrow \text{eval}_{\text{PRF}}(\bar{\kappa}, \text{key expansion}|\eta_C|\eta_{sid})$.
6. Generate confirmation messages $(\xi_C, \delta_C) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished}|h)$ and $(\xi_S, \delta_S) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished}|h)$.
7. Compute $\bar{\xi}_C$ and $\bar{\xi}_S$ by protecting ξ_C and ξ_S using the keys $\hat{\kappa}_C, \hat{\kappa}_S, \hat{\kappa}_{C,IV}, \hat{\kappa}_{S,IV}$ and authenticated encryption as specified by the cipher suite.
8. Output $(\kappa_C, \kappa_S, \kappa_{C,IV}, \kappa_{S,IV})$ and $(\bar{\xi}_C, \bar{\xi}_S)$.

The construction statement is almost the same as in Section 4; the only difference is that the keys are computed differently.

We aim for the statement that the protocol $(\text{expc13}, \text{exps13})$ constructs from $\text{MSK}_{N,\rho,AUX,n}$ the parallel composition of one copy of $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$ for each client $C \in \mathcal{C}$, each such resource corresponding to a server's sub-interface S/η for the nonce $\eta = \rho(C)$ chosen by client C . This, however, is not directly true since by the definition of the server's converter exps13 , an attacker can generate output at a server's sub-interface S/η for client nonces $\eta \notin \rho(\mathcal{C})$. This is not an artifice of our constructive treatment, but an inherent trait of protocols with unilateral authentication: the key may be shared either between clients and servers *or* between attackers and servers. We capture this feature at the level of the key expansion step by introducing "residual" resources $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$ for which the client's interface is also provided as a sub-interface of the E -interface. We denote the parallel composition of these resources as $\bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)}$, where the double-bracket notation indicates that the interfaces of each one of the (three-interface) resource $\stackrel{KSP,*}{=} \bullet_{\text{cphs},n}$ are provided as interfaces I and J for each $(I, J) \in \mathcal{P}$.

We show the following result, where we describe the pseudo-random function PRF by a converter prf and denote the system outputting a single 384-bit random string by \mathbf{U}_{384} and the uniform random function (with the same output length as PRF) by \mathbf{F} .

Lemma 17. *The protocol $(\text{expc13}, \text{exps13})$ constructs the keys $\bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)}$ from the handshake master secret $\text{MSK}_{N,\rho,AUX,n}$, for $\mathcal{P} = \{(C, S/\rho(C)) : C \in \mathcal{C}\} \cup \{(E/\eta, S/\eta) : \eta \notin \rho(\mathcal{C})\}$. The construction holds under the assumption that the hash function H is collision resistant and the pseudo-random function PRF is indistinguishable from a random function. In more detail, for the simulator σ and the reductions \mathbf{C} and \mathbf{C}' described in the proof,*

$$\text{MSK}_{N,\rho,AUX,n} \xrightarrow{(\text{expc13}, \text{exps13}), \sigma, (\varepsilon, \varepsilon)} \bigotimes_{(I,J) \in \mathcal{P}} \llbracket \stackrel{KSP,*}{=} \bullet_{\text{cphs},n} \rrbracket^{(I,J)}$$

such that for all distinguishers \mathbf{D} , $\varepsilon(\mathbf{D}) \preceq \Gamma^{\mathbf{DC}}(\mathbf{G}^{\mathbf{CR}}) + 2 \cdot |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$.

Proof sketch. The proof proceeds almost exactly as the one of Lemma 16; the difference is that the PRF is keyed with two different values per session (the handshake master secret and the main master secret, respectively), and this is accounted for by an additional factor 2 in the reduction. (The hybrid argument is performed over twice the number of intermediate steps.) On a very high level: by the collision resistance of the hash function, and since the pairs of nonces are unique by assumption, the inputs to the PRFs are either guaranteed to be the same (at the client and the server), or guaranteed to be distinct (for "different" uses or sessions). Then, by the assumption that the initial keys are unique, the remainder is reduced on the pseudo-randomness of the assumed PRF. \square

5 Constructing a Unilaterally Secure Channel

We describe the goal of the TLS *record layer* as constructing, from a unilateral key and insecure communication channels, a bidirectional unilaterally secure communication channel. TLS specifies several alternative cipher suites that are supposed to achieve this constructive step. The most widely used ones are based on an Authenticate-then-Encrypt combination of a MAC scheme and a symmetric encryption scheme, but [DR08] also specifies the possibility of using a monolithic authenticated encryption scheme. In this section, we prove the Authenticate-then-Encrypt modes based on [MT10], leaving the other cipher suites for future work.

As discussed in Section 1 and in previous work, the “finished” message of the TLS protocol cannot be regarded as part of the handshake if one wants to prove a strong security notion for the key. (The reason is that the actual key is used to protect the finished messages, which allows to verify whether an obtained key is correct.) Hence, as [JKSS12], we prove the record layer protocol *including* the finished messages and with the assumption of only a unilaterally authenticated key. The resource we want to construct by the record protocol is the “unilateral” channel $\leftarrow^* \rightarrow \bullet_n$ described below.

$\leftarrow^* \rightarrow \bullet_n$
<p>No attacker present: Behave as a (secure multi-message) channel between interfaces C and $S/1$.</p> <p>Attacker present:</p> <ul style="list-style-type: none"> • Upon the <i>first</i> input (key-c, e) with $e \in [n]$ at the E-interface (if e was not used before), provide a secure multiple-use (i.e., keep a buffer of undelivered messages) channel between C and S/e. In particular: <ul style="list-style-type: none"> – On input a message $m \in \mathcal{B}^{\leq 16384}$ at the C-interface, output m at interface E. – On input (deliver, client) at the E-interface, deliver the next message at S/e. – On input a message $m' \in \mathcal{B}^{\leq 16384}$ at the S/e-interface, output m' at interface E. – On input (deliver, server) at the E-interface, deliver the next message at C. • After input (conquer, e) with $e \in [n]$ at the E-interface (if e was not used before), forward messages in $\mathcal{B}^{\leq 16384}$ bidirectionally between the S/e- and E/e-interfaces.

Generally, the record layer protocol is a pair of converters which both obtain at the respective inside interfaces keys and finished messages (as given by the $\stackrel{KSP,*}{=} \bullet$ -resource). The client’s converter first sends the ξ_C -message (authenticated and encrypted), and then obtains, decrypts, and checks the ξ_S -message. If the check succeeds, payload messages are processed and transmitted. The server’s converter first waits for the (encrypted) ξ_C -message, decrypts, and checks. If successful, the converter sends ξ_S authenticated and encrypted, and later processes and transmits payload messages. More precisely, the two converters are described as follows.

The client’s converter behaves as follows.

1. Obtain at the inside interface keys $\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV}$ and “finished” messages ξ_C, ξ_S . Send the message ξ_C via the inside interface.
2. Upon receiving a message at the inside interface, compare it to ξ_S . If this step fails, abort.
3. Messages obtained at the outside interface are processed with the scheme(s) specified by the respective cipher suite using the keys $\kappa_{C,a}$ and $\kappa_{C,e}$ and sent via the inside interface. Further ciphertexts obtained at the inside interface are also processed with $\kappa_{S,e}$ and $\kappa_{S,a}$, the plaintexts are output at the outside interface. If any (MAC) verification fails, abort.

The server's converter behaves as follows:

1. Obtain at the inside interface keys $\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV}$ and “finished” messages ξ_C, ξ_S .
2. Upon receiving a message at the inside interface, compare it to ξ_C . If the above step fails, abort. Send the message ξ_S via the inside interface.
3. Messages obtained at the outside interface are processed with the scheme(s) (using the keys $\kappa_{S,a}$ and $\kappa_{S,e}$) and sent via the inside interface. Further ciphertexts obtained at the inside interface are processed (with $\kappa_{C,e}$ and $\kappa_{C,a}$), the plaintexts are output at the outside interface. If any (MAC) verification fails, abort.

5.1 Cipher Suites based on Stream Ciphers

The TLS standard [DR08] describes the record layer protocol based on stream ciphers in Section 6.2.3.1, the standard cipher suites using this type of encryption scheme are based on the cipher RC4. We prove the security following [MT10, Corollary 1] based on the assumption that the stream cipher produces a stream of pseudo-random bits.²⁰ In more detail, we formalize the assumption on a stream cipher by requiring that the distinguishing advantage between stream \mathbf{U}_k (i.e., the stream generated by the cipher when initialized with a uniformly random k -bit key) and \mathbf{U}^* (a stream of uniformly random bits²¹) is small. The scheme is formalized as the pair $(\text{atec}, \text{ates})$ of converters.

Lemma 18. *The protocol $(\text{atec}, \text{ates})$ constructs from $\stackrel{KSP,*}{=} \bullet$ the channel $\leftarrow^* \rightarrow \bullet_n$, under the assumptions that the used stream cipher is pseudo-random and HMAC is strongly unforgeable. More formally, for the simulator σ and the reductions \mathbf{C}, \mathbf{C}' described in the proof,*

$$\stackrel{KSP,*}{=} \bullet \quad \xRightarrow{(\text{atec}, \text{ates}), \sigma, (0, \varepsilon)} \quad \leftarrow^* \rightarrow \bullet_n,$$

with $\varepsilon(\mathbf{D}) \doteq 2 \cdot \Delta^{\mathbf{DC}}(\text{stream } \mathbf{U}_k, \mathbf{U}^*) + 2 \cdot \Gamma^{\mathbf{DC}'}(\mathbf{G}^{\text{suf-cma}})$ for all distinguishers \mathbf{D} .

Proof sketch. The validity of the availability condition follows because the resource $\stackrel{KSP,*}{=} \bullet$ outputs the same keys and “finished” messages at interfaces C and S , so the verification of these messages succeeds as atec and ates compute the same key streams and MACs.

To prove the validity of the security condition, we describe a simulator σ that initializes bits $b_e = 0$ for each $e \in [n]$. Then:

- Upon input $(\text{key-c}, e)$ at the outside interface, if $e \in [n]$ and $b_e = 0$, then set $b_e = 1$ and $\bar{e} = e$. Simulate the transmission of the client's finished message (a uniform random string of length 256 bits—96 bits “finished” message and 160 bits MAC) as the first message c_1 from C to S/e .
- Once both (deliver) has been input at the outside interface and the first message \tilde{c}_1 is delivered to S in session \bar{e} , if $c_1 = \tilde{c}_1$ then simulate a finished message from the server to the client, again by choosing a 256-bit string c_2 uniformly at random.

In the following, upon input the i th message length ℓ_i corresponding to a server message in session \bar{e} at the inside interface, output a uniformly random string of length $\ell_i + 160$. Also, whenever a message is delivered (via the outer interface) to the server session \bar{e} , check whether *exactly* the same messages were simulated as being sent by the client before. In

²⁰As demonstrated in [ABP⁺13], the assumption that RC4 is pseudo-random is dangerous. The proof holds for arbitrary stream ciphers.

²¹Formally, \mathbf{U}^* and stream allow to obtain the stream by querying for one bit at a time.

this case, input (ack, C) at the inside interface, otherwise halt the server session \bar{e} (i.e., stop processing messages for this server session, and, in real implementations, send an alert; however, for our treatment we omit alerts from the protocol description.).

- Once the first message \tilde{c}_2 is delivered to C , if $c_2 = \tilde{c}_2$ then record the client as active. In the following, upon input the i th message length ℓ_i corresponding to a client message at the inside interface, output a uniformly random string of length $\ell_i + 160$. Also, as above, whenever a message is delivered to the client, if *exactly* the same sequence of message was simulated as being sent by the server's session \bar{e} before, then input (ack, S) at the inside interface, otherwise halt the client (i.e., stop processing messages for the client—note that we simplify the protocol and do not handle error messages).
- Upon input $(\text{inject}, e, \bar{\kappa}_{C,a}, \bar{\kappa}_{C,e}, \bar{\kappa}_{S,a}, \bar{\kappa}_{S,e}, \bar{\kappa}_{C,IV}, \bar{\kappa}_{S,IV}, \bar{\xi}_C, \bar{\xi}_S)$ at the E -interface with $e \in [n]$ and $b_e = 0$, set $b_e = 1$ and record the values. When the first message is delivered to the session e , check whether the message is a correctly MAC'ed and encrypted version (with $\bar{\kappa}_{C,a}$ and $\bar{\kappa}_{C,e}$) of $\bar{\xi}_C$ (if not, abort the server session e). Respond with a correctly MAC'ed and encrypted version (with $\bar{\kappa}_{S,a}$ and $\bar{\kappa}_{S,e}$) of $\bar{\xi}_S$. Subsequently, MAC and encrypt messages sent in the server session e with the keys $\bar{\kappa}_{S,a}$ and $\bar{\kappa}_{S,e}$. For messages given at the outside interface for this session, decrypt with $\bar{\kappa}_{C,e}$ and verify the MAC with $\bar{\kappa}_{C,a}$. In case of success, inject the resulting message via the inside interface, otherwise halt the server session e .

First, we note that the simulation of all sessions except for \bar{e} is perfect, as the simulator makes exactly the same computations as the protocol.

To prove the security statement, we use a hybrid system \mathbf{H}_1 similarly to $\sigma^E \leftarrow^* \rightarrow_{\bullet n}$ with the difference that the key stream is generated by the stream cipher with a uniformly random key instead of uniformly at random. The reduction system \mathbf{C} simulates all sessions similarly to $\sigma^E \leftarrow^* \rightarrow_{\bullet n}$, but in session \bar{e} it uses the key stream from the connected system (with probability $\frac{1}{2}$ it does so for the client while using a fully random stream for the server, with the remaining probability it uses the given stream for the client and generates the server's stream using the stream cipher). This means that $\Delta^{\mathbf{D}} \left(\text{atec}^C \text{ates}^S \stackrel{KSP,*}{=} \bullet, \mathbf{H}_1 \right) \leq 2 \cdot \Delta^{\mathbf{DC}}(\text{stream } \mathbf{U}_{128}, \mathbf{U}^*)$.

Then, we use [MT10, Corollary 1] twice, once for each direction, to obtain the statement $\Delta^{\mathbf{D}} \left(\mathbf{H}_1, \sigma^E \leftarrow^* \rightarrow_{\bullet n} \right) \leq 2 \cdot \Gamma^{\mathbf{DC}'}(\mathbf{G}^{\text{suf-cma}})$ and apply the triangular inequality to conclude. \square

5.2 Cipher Suites based on CBC Encryption

The TLS standard [DR08] describes the record layer protocol based on CBC encryption in Section 6.2.3.2, the standard cipher suites using this encryption mode are based on either 3DES or AES. We prove the security following [MT10, Corollary 2], based on the assumption that the used block cipher is a (super²²) PRP. In more detail, we formalize the assumptions on the block ciphers by requiring that the distinguishing advantage between $\text{bc } \mathbf{U}_k$ (i.e., the block cipher bc with block length ℓ which might e.g. be 3DES or AES initialized with a uniformly random key of appropriate length) and \mathbf{P}_ℓ (a uniformly random permutation) is small, even allowing both forward and backward queries.

In the following lemma, we use the converters atec' and ates' that implement the Authenticate-then-Encrypt composition of a CBC-mode with block cipher bc and a strongly unforgeable MAC. As we base the proof on [MT10, Corollary 2], it only applies to the case where the padding used by TLS is unique in the sense that it is the shortest possible such padding (and no length-hiding techniques are used).

²²In the reduction, it is necessary to make *inverse* queries to the permutation. This is unclear in [MT10].

Lemma 19. *The protocol $(\text{atec}', \text{ates}')$ constructs from $\stackrel{KSP,*}{=} \bullet$ the channel $\stackrel{*}{\leftarrow} \bullet_n$, under the assumptions that bc is a (super) PRP and HMAC is strongly unforgeable. More formally, for the simulator σ and the reductions \mathbf{C}, \mathbf{C}' described in the proof,*

$$\stackrel{KSP,*}{=} \bullet \quad \xRightarrow{(\text{atec}', \text{ates}'), \sigma, (0, \varepsilon)} \quad \stackrel{*}{\leftarrow} \bullet_n,$$

with $\varepsilon(\mathbf{D}) = 2 \cdot \Delta^{\text{DC}}(\text{bc } \mathbf{U}_k, \mathbf{P}_\ell) + 2 \cdot \Gamma^{\text{DC}'}(\mathbf{G}^{\text{suf-cma}}) + \frac{(ql)^2}{2^{\ell-1}}$ for all distinguishers \mathbf{D} .

Proof sketch. The availability condition follows as in Lemma 18.

To prove the validity of the security condition, we describe a simulator σ that initializes bits $b_e = 0$ for each $e \in [n]$. Then:

- Upon input $(\text{key-c}, e)$ at the outside interface, if $e \in [n]$ and $b_e = 0$, then set $b_e = 1$ and $\bar{e} = e$. Simulate the transmission of the client's finished message (a uniform random string of length $\ell \cdot \lceil \frac{256}{\ell} \rceil$ bits—96 bits “finished” message and 160 bits MAC, padded to the next block size) as the first message c_1 from C to S/e .
- Once both (deliver) has been input at the outside interface and the first message \tilde{c}_1 is delivered to S in session \bar{e} , if $c_1 = \tilde{c}_1$ then simulate a finished message from the server to the client, again by choosing a bit string c_2 of appropriate length uniformly at random. In the following, upon input the i th message length ℓ_i corresponding to a server message in session \bar{e} at the inside interface, output a uniformly random string of length $\ell \cdot \lceil \frac{\ell_i + 160}{\ell} \rceil$. Whenever a message is delivered (via the outer interface) to the server session \bar{e} , behave as the simulator in Lemma 18.
- Once the first message \tilde{c}_2 is delivered to C , if $c_2 = \tilde{c}_2$ then record the client as active. In the following, upon input the i th message length ℓ_i corresponding to a client message at the inside interface, output a uniformly random string of length $\ell \cdot \lceil \frac{\ell_i + 160}{\ell} \rceil$. Also, as above, whenever a message is delivered to the client, behave as the simulator in Lemma 18.
- Upon input $(\text{inject}, e, \bar{\kappa}_{C,a}, \bar{\kappa}_{C,e}, \bar{\kappa}_{S,a}, \bar{\kappa}_{S,e}, \bar{\kappa}_{C,IV}, \bar{\kappa}_{S,IV}, \bar{\xi}_C, \bar{\xi}_S)$ at the E -interface with $e \in [n]$ and $b_e = 0$, set $b_e = 1$ and record the values. When the first message is delivered to the session e , check whether the message is a correctly MAC'ed, padded, and encrypted version (with $\bar{\kappa}_{C,a}$ and $\bar{\kappa}_{C,e}$) of $\bar{\xi}_C$ (if not, abort the server session e). Respond with a correctly MAC'ed, padded, and encrypted version (with $\bar{\kappa}_{S,a}$ and $\bar{\kappa}_{S,e}$) of $\bar{\xi}_S$. Subsequently, MAC, pad, and encrypt messages sent in the server session e with the keys $\bar{\kappa}_{S,a}$ and $\bar{\kappa}_{S,e}$. For messages given at the outside interface for this session, decrypt with $\bar{\kappa}_{C,e}$ and verify the padding and the MAC with $\bar{\kappa}_{C,a}$. In case of success, inject the resulting message via the inside interface, otherwise halt the server session e .

First, we note that the simulation of all sessions except for \bar{e} is perfect, as the simulator makes exactly the same computations as the protocol.

To prove the security statement, we use a hybrid system \mathbf{H}_1 similarly to $\sigma^E \stackrel{*}{\leftarrow} \bullet_n$ with the difference that the CBC scheme is computed using a uniformly random permutation \mathbf{P}_ℓ (instead of bc). The reduction system \mathbf{C} simulates all sessions similarly to $\sigma^E \stackrel{*}{\leftarrow} \bullet_n$, but in session \bar{e} it uses the connected system (with probability $\frac{1}{2}$ it does so for the client while using a fully random permutation for the server, with the remaining probability it uses the given permutation for the server and generates the server's stream using bc). This means that $\Delta^{\mathbf{D}} \left(\text{atec}'^C \text{ates}'^S \stackrel{KSP,*}{=} \bullet, \mathbf{H}_1 \right) \leq 2 \cdot \Delta^{\text{DC}}(\text{bc } \mathbf{U}_k, \mathbf{P}_\ell)$. Then, we use [MT10, Corollary 2] twice, once for each direction, to obtain the statement $\Delta^{\mathbf{D}} \left(\mathbf{H}_1, \sigma^E \stackrel{*}{\leftarrow} \bullet_n \right) \leq 2 \cdot \Gamma^{\text{DC}'}(\mathbf{G}^{\text{suf-cma}}) + \frac{(ql)^2}{2^{\ell-1}}$ and apply the triangular inequality to conclude. \square

5.3 Cipher Suites based on AEAD Encryption

The current draft of TLS 1.3 [DR15] describes (only) a record layer protocol which is based on authenticated encryption with associated data (AEAD). This mode has been analyzed by Badertscher et al. [BMM⁺15] in recent work. Their result can be “imported” into our work along the lines of Sections 5.1 and 5.2.

We prove the security following [BMM⁺15, Theorems 1 and 2], based on the assumption that the used authenticated encryption is secure. As described in Appendix C.5, we formalize the security of authenticated encryption via the distinguishing advantage between the two systems $\mathbf{G}_0^{\text{aead}}$ and $\mathbf{G}_1^{\text{aead}}$.

In the following lemma, we use the converters `aeadc` and `aeads` that use AEAD encryption as described in [DR15] with sequence numbers as nonces. We base the proof on [BMM⁺15, Theorems 1 and 2].

Lemma 20. *The protocol $(\text{aeadc}, \text{aeads})$ constructs from $\stackrel{KSP,*}{=} \bullet$ the channel $\leftarrow^* \rightarrow_{\bullet n}$, under the assumption that the underlying AEAD cipher is secure. More formally, for the simulator σ and the reduction \mathbf{C} described in the proof,*

$$\stackrel{KSP,*}{=} \bullet \quad \xRightarrow{(\text{aeadc}, \text{aeads}), \sigma, (0, \varepsilon)} \quad \leftarrow^* \rightarrow_{\bullet n},$$

with $\varepsilon(\mathbf{D}) = 2 \cdot \Delta^{\mathbf{DC}}(\mathbf{G}_0^{\text{aead}}, \mathbf{G}_1^{\text{aead}})$ for all distinguishers \mathbf{D} .

Proof sketch. The availability condition follows as in Lemma 18.

To prove the validity of the security condition, we describe a simulator σ that initializes bits $b_e = 0$ for each $e \in [n]$. Then:

- Upon input $(\text{key-c}, e)$ at the outside interface, if $e \in [n]$ and $b_e = 0$, then set $b_e = 1$ and $\bar{e} = e$. Simulate the transmission of the client’s finished message (a uniform random string of length $\ell(96)$ bits—96 bits “finished” message encrypted with AEAD cipher) as the first message c_1 from C to S/e .
- Once both (deliver) has been input at the outside interface and the first message \tilde{c}_1 is delivered to S in session \bar{e} , if $c_1 = \tilde{c}_1$ then simulate a finished message from the server to the client, again by choosing a bit string c_2 of appropriate length uniformly at random. In the following, upon input the i th message length ℓ_i corresponding to a server message in session \bar{e} at the inside interface, output a uniformly random string of length $\ell(\ell_i)$. Whenever a message is delivered (via the outer interface) to the server session \bar{e} , behave as the simulator in Lemma 18.
- Once the first message \tilde{c}_2 is delivered to C , if $c_2 = \tilde{c}_2$ then record the client as active. In the following, upon input the i th message length ℓ_i corresponding to a client message at the inside interface, output a uniformly random string of length $\ell(\ell_i)$. Also, as above, whenever a message is delivered to the client, behave as the simulator in Lemma 18.
- Upon input $(\text{inject}, e, \bar{\kappa}_C, \bar{\kappa}_S, \bar{\kappa}_{C,IV}, \bar{\kappa}_{S,IV}, \bar{\xi}_C, \bar{\xi}_S)$ at the E -interface with $e \in [n]$ and $b_e = 0$, set $b_e = 1$ and record the values. When the first message is delivered to the session e , check whether the message is a correct according to the AEAD scheme (with $\bar{\kappa}_C$) of $\bar{\xi}_C$ (if not, abort the server session e). Respond with a correctly AEAD encrypted version (with $\bar{\kappa}_S$) of $\bar{\xi}_S$. Subsequently, AEAD encrypt messages sent in the server session e with the key $\bar{\kappa}_S$. For messages given at the outside interface for this session, decrypt with $\bar{\kappa}_C$. In case of success, inject the resulting message via the inside interface, otherwise halt the server session e .

First, we note that the simulation of all sessions except for \bar{e} is perfect, as the simulator makes exactly the same computations as the protocol.

To prove the security statement, we use [BMM⁺15, Theorems 1 and 2] twice, once for each direction, to obtain the statement $\Delta^D \left(\text{aeadc}^A \text{aeads}^B, \sigma^E \xleftarrow{*} \xrightarrow{\bullet_n} \right) \leq 2 \cdot \Delta^{DC'} (\mathbf{G}_0^{\text{aead}}, \mathbf{G}_1^{\text{aead}})$ via the triangular inequality. \square

6 Reconstructing TLS

In this section we argue that the composition of the converters we presented in Sections 3.3.2, 4, and 5 forms in fact the TLS protocol, and then give the full security statements for each of the cipher suites.

Note that in our deconstruction of TLS, we use an intermediate converter that does not appear in the TLS-DH and TLS-RSA versions, namely constructing the authenticated network resource $\succsim \bullet_{N,\rho,\mathfrak{F},\text{SIG},n}$. As the difference between the TLS-DH/TLS-RSA and the TLS-DHE protocols only appears in the construction of the master secret resource (in our de-construction, obtaining the master key resource $\text{MSK}_{N,\rho,AUX,n}$ from the nonce-exchange resource $\text{SNET}_{N,\rho,n}$), we only consider the TLS-DHE variant, noting that the same considerations hold for the TLS-DH/TLS-RSA versions.

In fact, we describe TLS as a composition of a client TLS-DHE converter (denoted as `tlsdhec`) and one server TLS-DHE converter. The client TLS-DHE converter `tlsdhec` represents the composition of all the converters mounted at the client interfaces of our respective resources, i.e., `tlsdhec` = `atec` \circ `expc` \circ `dhec` \circ `vrf` \circ `hec`. We explicitly outline the resulting `tlsdhec` converter next:

1. **hec**: Obtain a (random) nonce $\eta_C \in \mathcal{N}$ (at the inside interface) and send it to the server via the inside interface. Upon receiving nonce $\tilde{\eta}$ at the inside interface, output $(\eta_C, \tilde{\eta})$ at the outside interface.
2. **vrf**: Upon receiving message *cert* at the inside interface corresponding to $\text{SNET}_{N,\rho,n}$, query (`verify`, *cert*) at the inside sub-interface corresponding to $\text{PKI}_{\mathfrak{F}}$; abort if the verification fails or if *cert* is not a well-formed certificate $\text{cert} = (vk, f(vk))$. Upon obtaining a second message *m'* from $\text{SNET}_{N,\rho,n}$, parse *m'* as (m, s) (abort if that is impossible). If $\text{vrf}(\eta_C | \eta_{sid} | m, s; vk) = 1$, then output $(\text{cert}, \eta_C, \eta_{sid}, m, s)$ at the outside interface. (Otherwise abort.)
3. **dhec**: Parse message *m* (obtained at the inside interface) as $p|g|g'|s$ (abort if impossible). Choose $u \leftarrow \{1, \dots, q\}$ (with $q = |\mathbf{Z}_p^\times|$) and input g^u at the inside interface. Query $g'^u | \text{master secret} | \eta_C | \eta_{sid}$ at RO_{384} , in order to obtain a key $\kappa \in \{0, 1\}^{384}$. Output $(\kappa, \eta_C, \eta_{sid}, aux | m | s | g^u)$.
4. **expc**: Use the value κ to generate keys

$$(\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV}) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion} | \eta_C | \eta_{sid}).$$

Using the concatenation of the previously transmitted messages $m \doteq \text{cert} | p | g | g' | s$, compute the “finished” messages $(\xi_C, \delta_C) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished} | H(\eta_C | \eta_{sid} | m | \gamma))$ and $(\xi_S, \delta_S) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished} | H(\eta_C | \eta_{sid} | m | c_{\xi_C} | \gamma))$. In the above computation, the constant γ stands for the “ChangeCipherSpec” message, whereas the value c_{ξ_C} is computed as a function of ξ_C depending on the adopted cipher suite (in particular, ξ_C needs to be included in the hash of the message sent by the server). Output $(\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV})$ and (ξ_C, ξ_S) .

5. **atec**: Use the record layer scheme(s) specified by the cipher suite with the keys $\kappa_{C,a}$ and $\kappa_{C,e}$ obtained at the inside interface, respectively, to process the message ξ_C (obtained with the keys), and to computed and send the obtained ciphertext via the inside interface.

Upon receiving a message at the inside interface, process it using the keys $\kappa_{S,e}$ and $\kappa_{S,a}$. Compare the plaintext to ξ_S . If any one of the above steps fails, abort. From this point on, messages obtained at the outside interface are processed with the specified scheme(s) using keys $\kappa_{C,a}$ and $\kappa_{C,e}$ and sent via the inside interface. Further ciphertexts obtained at the inside interface are processed with the keys $\kappa_{S,e}$ and $\kappa_{S,a}$, the plaintexts are output at the outside interface. If any (MAC) verification fails, abort.

Note that the composite converter corresponds exactly to the client's protocol in TLS, with the specification that, in the expansion step, the client first computes ξ_S as the PRF, under the key κ of the hash of all the past messages, including ξ_C . We note that our description of the protocol omits several constants that appear in the original protocol.

The server converter **tlsdhes** connects to the S -interface, and is composed of all the converters connected to the S -interface, i.e., $\text{tlsdhes} = \text{ates} \circ \text{exps} \circ \text{dhes}_{\mathcal{G}} \circ \text{sgn} \circ \text{hes}$. We explicitly outline the resulting **tlsdhes** converter below:

1. **hes**: Upon receiving a nonce $\tilde{\eta}_C$ at the inside C -sub-interface for some $C \in \mathcal{A}_{\text{TCP}}$, choose a nonce $\eta_{sid} \leftarrow \mathbb{N}$, send η_{sid} via the inside C -sub-interface. Output η_{sid} at outside sub-interface \tilde{sid} .
2. **sgn**: Initially compute $(sk, vk) \leftarrow \text{gen}$. Input vk at the inside sub-interface corresponding to $\text{PKI}_{\mathfrak{F}}$, obtaining a response s , and set $\text{cert} = (vk, s)$. Output cert at the outside interface (as auxiliary information). Subsequently, for each inside sub-interface $sid = (\eta, e)$ outputting a nonce output η_{sid} , output η_{sid} at the respective outside sub-interface and send cert via $\text{SNET}_{\mathbb{N}, \rho, n}$ (in the respective session). Obtaining a message m at the outside (sub-interface for session sid), compute $s \leftarrow \text{sign}(\eta_C | \eta_{sid} | m, sk)$ and send (m, s) in session sid . Output s at the respective outside sub-interface.
3. **dhes $_{\mathcal{G}}$** : Upon obtaining a nonce η_{sid} at the inside sub-interface $sid = (\eta_C, e)$, choose a modulus $p \in \mathbb{N}$ and a generator $g \in \mathbf{Z}_p^\times$ according to the distribution \mathcal{G} . Also, choose an exponent $v \leftarrow \mathbb{N}$. Input the value $m = p|g|g^v$ at the respective inside sub-interface (obtaining a signature s in response). Upon receiving a group element \tilde{g} at the respective inside sub-interface, query $\tilde{g}^v | \text{master secret} | \eta_C | \eta_{sid}$ at RO_{384} , call the result κ . Output $(\kappa, \eta_{sid}, \text{aux} | m | s | \tilde{g})$ at the respective outside sub-interface.
4. **exps**: Obtaining $(\kappa, \eta_{sid}, \text{aux} | m | s | \tilde{g})$ at the inside sub-interface of session $sid = (\eta_C, e)$, generate keys $\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV} \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{key expansion} | \eta_C | \eta_{sid})$. Denote $m = \text{cert} | p | g | g^v | s$. Generate messages $(\xi_C, \delta_C) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{client finished} | H(\eta_C | \eta_{sid} | m | \gamma))$ and $(\xi_S, \delta_S) \leftarrow \text{eval}_{\text{PRF}}(\kappa, \text{server finished} | H(\eta_C | \eta_{sid} | m | c_{\xi_C} | \gamma))$. Output $(\kappa_{C,a}, \kappa_{C,e}, \kappa_{S,a}, \kappa_{S,e}, \kappa_{C,IV}, \kappa_{S,IV})$ and (ξ_C, ξ_S) at the outside sub-interface sid .
5. **ates**: There is one such converter for each nonce $\eta \in \mathcal{N}$, which connects at the respective sub-interface. Upon receiving a message at the inside interface, process the message with the scheme(s) specified by the cipher suite using the keys $\kappa_{C,e}$ and $\kappa_{C,a}$. Compare the plaintext to ξ_C . If any one of the above steps fails, abort. Use again the specified schemes, now using the keys $\kappa_{S,a}$ and $\kappa_{S,e}$, to process the message ξ_S , and send the obtained ciphertext via the inside interface. Messages obtained at the outside interface are processed with the specified scheme(s) using the keys $\kappa_{S,a}$ and $\kappa_{S,e}$ and sent via the inside interface. Further ciphertexts obtained at the inside interface are processed using $\kappa_{C,e}$ and $\kappa_{C,a}$, the plaintexts are output at the outside interface. If any (MAC) verification fails, abort.

Note once more that this amounts to the server protocol in TLS-DHE. In the following section we give the full security statements for all versions TLS-DH, TLS-DHE, and TLS-RSA.

6.1 Full Security Statements

Summarizing the bounds obtained in the previous sections, the full security statements for TLS-DH, TLS-DHE, and TLS-RSA are as follows. We write the theorems for the cipher suites based on the stream cipher only, but the analogous theorems for the CBC-based cipher suites are obtained in the same way.

We start by showing the complete security statement for a cipher suite based on TLS-DH and a stream cipher.

Theorem 21. *Let $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ be a set of clients. The TLS-DH protocol constructs, for each client $C \in \mathcal{C}$, one unilaterally secure channel $\leftarrow^* \rightarrow_{\bullet n}$ from NET , PKI , and RO_{384} . Concretely, for the simulator σ and the reductions \mathbf{C} , \mathbf{C}' , \mathbf{C}'' , \mathbf{C}''' , $\mathbf{C}^{(iv)}$, $\mathbf{C}^{(v)}$, and $\mathbf{C}^{(vi)}$ obtained by composing those described in the lemmas:*

$$[\text{NET}, \text{PKI}, \text{RO}_{384}] \xRightarrow{(\text{tlsdhec}, \text{tlsdhes}), \sigma, (\varepsilon_1, \varepsilon_2)} \bigotimes_{(I, J) \in \mathcal{P}} \llbracket \leftarrow^* \rightarrow_{\bullet n} \rrbracket^{(I, J)},$$

with $\varepsilon_1(\mathbf{D}) = \binom{|\mathcal{C}|}{2} \cdot 2^{-224} + \Gamma^{\text{DC}}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\text{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$ and

$$\begin{aligned} \varepsilon_2(\mathbf{D}) = & \left(\binom{n}{2} + \binom{|\mathcal{C}|}{2} \right) \cdot 2^{-224} + \Gamma^{\text{DC}''}(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}) + \Gamma^{\text{DC}'''}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\text{DC}^{(iv)}}(\text{prf } \mathbf{U}_{384}, \mathbf{F}) \\ & + |\mathcal{C}| \left(2 \cdot \Delta^{\text{DC}^{(v)}}(\text{stream } \mathbf{U}_k, \mathbf{U}^*) + 2 \cdot \Gamma^{\text{DC}^{(vi)}}(\mathbf{G}^{\text{suf-cma}}) \right), \end{aligned}$$

for all distinguishers \mathbf{D} and with the set \mathcal{P} defined as in Lemma 16.

Proof. For the availability condition, we lose a term $\binom{|\mathcal{C}|}{2} 2^{-224}$ in the construction of the resource NAME_{ρ} from scratch (see Section 3.2.1), and another term $\Gamma^{\text{DC}}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\text{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$ in the construction of $\left[\bigotimes_{C \in \mathcal{C}} \llbracket \xRightarrow{KSP, *} \rrbracket_{\text{cphs}, n}^{(C, S/\rho(C))}, \bigotimes_{\eta \in \mathcal{N} \setminus \rho(\mathcal{C})} \llbracket \tilde{\mathbf{R}}_{\text{cphs}, n} \rrbracket^{(S/\eta)} \right]$ via $(\text{expc}, \text{exps})$ (see Lemma 16). The converters ates connected to the interfaces corresponding to $\eta \in \mathcal{N} \setminus \rho(\mathcal{C})$ do not obtain keys and hence remain inactive.

For the security condition, we again lose a term $\binom{|\mathcal{C}|}{2} 2^{-224}$ when constructing the resource NAME_{ρ} from scratch. We also lose: a term $\binom{n}{2} 2^{-224}$ in the construction of $\text{SNET}_{\mathbf{N}, \rho, n}$ from the resources NET and NAME_{ρ} (see Lemma 8); a term $\Gamma^{\text{DC}}(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}})$ in the construction of the resource $\text{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ from the resource $[\text{RO}_{384}, \text{SNET}_{\mathbf{N}, \rho, n}, \text{PKI}_{\mathcal{F}}]$ (see Lemma 9); a term $\Gamma^{\text{DC}'}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\text{DC}''}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$ in constructing the parallel composition of keys $\left[\bigotimes_{C \in \mathcal{C}} \llbracket \xRightarrow{KSP, *} \rrbracket_{\text{cphs}, n}^{(C, S/\rho(C))}, \bigotimes_{\eta \in \mathcal{N} \setminus \rho(\mathcal{C})} \llbracket \tilde{\mathbf{R}}_{\text{cphs}, n} \rrbracket^{(S/\eta)} \right]$ via $(\text{expc}, \text{exps})$ (see Lemma 16); finally, noting that we lose a term $2 \cdot \Delta^{\text{DC}'''}(\text{stream } \mathbf{U}_k, \mathbf{U}^*) + 2 \cdot \Gamma^{\text{DC}^{(iv)}}(\mathbf{G}^{\text{suf-cma}})$ for each obtained channel between a client and server. The $|\mathcal{C}|$ -many such channels are obtained by parallel composition. \square

We obtain the following analogous statement for a cipher suite based on TLS-DHE together with a stream cipher.

Theorem 22. *Let $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ be a set of clients. The TLS-DHE protocol constructs, for each client $C \in \mathcal{C}$, one unilaterally secure channel $\leftarrow^* \rightarrow_{\bullet n}$ from NET , PKI , and RO_{384} . Concretely, for the simulator σ and the reductions \mathbf{C} , \mathbf{C}' , \mathbf{C}'' , \mathbf{C}''' , $\mathbf{C}^{(iv)}$, $\mathbf{C}^{(v)}$, $\mathbf{C}^{(vi)}$, and $\mathbf{C}^{(vii)}$ obtained by composing the reductions from the lemmas,*

$$[\text{NET}, \text{PKI}, \text{RO}_{384}] \xRightarrow{(\text{tlsdhec}, \text{tlsdhes}), \sigma, (\varepsilon_1, \varepsilon_2)} \bigotimes_{(I, J) \in \mathcal{P}} \llbracket \leftarrow^* \rightarrow_{\bullet n} \rrbracket^{(I, J)},$$

with $\varepsilon_1(\mathbf{D}) = \binom{|\mathcal{C}|}{2} \cdot 2^{-224} + \Gamma^{\mathbf{DC}}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$ and

$$\begin{aligned} \varepsilon_2(\mathbf{D}) = & 2 \cdot \left(\binom{n}{2} + \binom{|\mathcal{C}|}{2} \right) \cdot 2^{-224} + \Gamma^{\mathbf{DC}''}(\mathbf{G}^{\text{uf-cma}}) + n \cdot |\mathcal{C}| \cdot \Gamma^{\mathbf{DC}'''}(\mathbf{G}_g^{\text{GapDH}}) \\ & + \Gamma^{\mathbf{DC}^{(iv)}}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\mathbf{DC}^{(v)}}(\text{prf } \mathbf{U}_{384}, \mathbf{F}) \\ & + |\mathcal{C}| \cdot \left(2 \cdot \Delta^{\mathbf{DC}^{(vi)}}(\text{stream } \mathbf{U}_k, \mathbf{U}^*) + 2 \cdot \Gamma^{\mathbf{DC}^{(vii)}}(\mathbf{G}^{\text{suf-cma}}) \right), \end{aligned}$$

for all distinguishers \mathbf{D} and with the set \mathcal{P} defined as in Lemma 16.

Proof. The proof of the availability condition follows exactly the same scheme as above. A similar argument holds for the security condition, with the exception that we lose a term $\Gamma^{\mathbf{DC}}(\mathbf{G}^{\text{uf-cma}})$ while constructing the resource $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$ from the resources $\mathbf{SNET}_{\mathbf{N}, \rho, n}$ and $\mathbf{PKI}_{\mathfrak{F}}$ (see Lemma 10), and then a term $n \cdot |\mathcal{C}| \cdot \Gamma^{\mathbf{DC}}(\mathbf{G}_g^{\text{GapDH}})$ in constructing the resource $\mathbf{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ from the resources \mathbf{RO}_{384} and $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$ (see Lemma 11). The remainder of the proof is as above. \square

The analogous result also holds with respect to a cipher suite based on TLS-RSA together with a stream cipher. Note that the entire security statement is implicitly parametrized by *gen*, the RSA key generation algorithm, which is specialized for the considered key length. In particular, this affects the converters *rsas* and *tlrsas* as well as the game $\mathbf{G}^{\text{nr-pca}}$.

Theorem 23 (TLS-RSA). *Let $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ be a set of clients. The TLS-RSA protocol constructs, for each client $C \in \mathcal{C}$, one unilaterally secure channel $\leftarrow^* \rightarrow_{\bullet n}$ from \mathbf{NET} , \mathbf{PKI} , and \mathbf{RO}_{384} . Concretely, there are a simulator σ and reductions \mathbf{C} , \mathbf{C}' , \mathbf{C}''_q , for $q \in \mathbb{N}$, \mathbf{C}''' , $\mathbf{C}^{(iv)}$, $\mathbf{C}^{(v)}$, and $\mathbf{C}^{(vi)}$ such that:*

$$[\mathbf{NET}, \mathbf{PKI}, \mathbf{RO}_{384}] \xRightarrow{(\text{tlrsac}, \text{tlrsas}), \sigma, (\varepsilon_1, \varepsilon_2)} \bigotimes_{(I, J) \in \mathcal{P}} \llbracket \leftarrow^* \rightarrow_{\bullet n} \rrbracket^{(I, J)},$$

with \mathcal{P} as defined in Lemma 16, $\varepsilon_1(\mathbf{D}) = \binom{|\mathcal{C}|}{2} \cdot 2^{-224} + \Gamma^{\mathbf{DC}}(\mathbf{G}^{\text{CR}}) + |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$ and

$$\begin{aligned} \varepsilon_2(\mathbf{D}) = & \left(\binom{n}{2} + \binom{|\mathcal{C}|}{2} \right) \cdot 2^{-224} + n \cdot |\mathcal{C}| \cdot \Gamma_q^{\mathbf{DC}''_q}(\mathbf{G}^{\text{nr-pca}}) + \frac{q}{2^{368}} + \Gamma^{\mathbf{DC}'''}(\mathbf{G}^{\text{CR}}) \\ & + |\mathcal{C}| \cdot \Delta^{\mathbf{DC}^{(iv)}}(\text{prf } \mathbf{U}_{384}, \mathbf{F}) + |\mathcal{C}| \cdot \left(2 \cdot \Delta^{\mathbf{DC}^{(v)}}(\text{stream } \mathbf{U}_k, \mathbf{U}^*) + 2 \cdot \Gamma^{\mathbf{DC}^{(vi)}}(\mathbf{G}^{\text{suf-cma}}) \right). \end{aligned}$$

Proof. The proof of the availability condition follows the same scheme as for DH. A similar argument holds for the security condition, except that we use the bounds for RSA for constructing the resource $\mathbf{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ (see Lemma 12). \square

Finally, we combine the results for TLS 1.3 in the same way as above.

Theorem 24. *Let $\mathcal{C} \subseteq \mathcal{A}_{\text{TCP}}$ be a set of clients. The TLS 1.3 protocol constructs, for each client $C \in \mathcal{C}$, one unilaterally secure channel $\leftarrow^* \rightarrow_{\bullet n}$ from \mathbf{NET} and \mathbf{PKI} . Concretely, for the simulator σ and the reductions \mathbf{C} , \mathbf{C}' , \mathbf{C}'' , \mathbf{C}''' , $\mathbf{C}^{(iv)}$, $\mathbf{C}^{(v)}$, $\mathbf{C}^{(vi)}$, and $\mathbf{C}^{(vii)}$ obtained by composing the reductions from the lemmas,*

$$[\mathbf{NET}, \mathbf{PKI}] \xRightarrow{(\text{tls13c}, \text{tls13s}), \sigma, (\varepsilon_1, \varepsilon_2)} \bigotimes_{(I, J) \in \mathcal{P}} \llbracket \leftarrow^* \rightarrow_{\bullet n} \rrbracket^{(I, J)},$$

with $\varepsilon_1(\mathbf{D}) = \binom{|\mathcal{C}|}{2} \cdot 2^{-224} + \Gamma^{\mathbf{DC}}(\mathbf{G}^{\mathbf{CR}}) + 2 \cdot |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'}(\text{prf } \mathbf{U}_{384}, \mathbf{F})$ and

$$\begin{aligned} \varepsilon_2(\mathbf{D}) = & 2 \cdot \left(\binom{n}{2} + \binom{|\mathcal{C}|}{2} \right) \cdot 2^{-224} + \Gamma^{\mathbf{DC}''}(\mathbf{G}^{\text{uf-cma}}) + n \cdot |\mathcal{C}| \cdot \Delta^{\mathbf{DC}'''}((g^A, g^B, g^{AB}), (g^A, g^B, g^C)) \\ & + 3 \cdot \Gamma^{\mathbf{DC}^{(iv)}}(\mathbf{G}^{\mathbf{CR}}) + 4 \cdot |\mathcal{C}| \cdot \Delta^{\mathbf{DC}^{(v)}}(\text{prf } \mathbf{U}_{384}, \mathbf{F}) \\ & + 2 \cdot |\mathcal{C}| \cdot \Delta^{\mathbf{DC}^{(vi)}}(\mathbf{G}_0^{\text{aead}}, \mathbf{G}_1^{\text{aead}}), \end{aligned}$$

for all distinguishers \mathbf{D} and with the set \mathcal{P} defined as in Lemma 16.

Proof. The proof of the availability condition follows exactly the same scheme as above. A similar argument holds for the security condition, with the exception that we lose a term $\Gamma^{\mathbf{DC}}(\mathbf{G}^{\text{uf-cma}})$ while constructing the resource $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$ from the resources $\mathbf{SNET}_{\mathbf{N}, \rho, n}$ and $\mathbf{PKI}_{\mathfrak{F}}$ (see Lemma 10), and then a term $n \cdot |\mathcal{C}| \cdot \Gamma^{\mathbf{DC}}(\mathbf{G}_G^{\text{GapDH}})$ in constructing the resource $\mathbf{MSK}_{\mathbf{N}, \rho, \text{AUX}, n}$ from the resources \mathbf{RO}_{384} and $\succsim \bullet_{\mathbf{N}, \rho, \mathfrak{F}, \text{SIG}, n}$ (see Lemma 11). The remainder of the proof is as above. \square

7 Conclusion and Lessons Learned

TLS can be seen as a composition of sub-protocols, each *constructing* a resource transmitting the remaining TLS messages as its payload. This approach allows to prove TLS in a modular way. In particular, we analyze all three unilateral key-exchange modes of TLS 1.2 in isolation, independently of the remainder of the protocol. We also analyze the the handshake of the recent TLS 1.3 draft. Interestingly, although the protocol has been changed considerably, we can use the same decomposition and only need to re-prove the protocol steps that were changed between the version. In both cases of TLS 1.2 and 1.3, we can (essentially) re-use statements on the record layer protocol proven in previous work, and obtain the security of the full protocol by the composition theorem of constructive cryptography. This approach, however, comes at the expense of resource complexity and strong assumptions (random oracles for all of TLS 1.2 and NR-PCA instead of OW-PCA for the RSA-PKCS handshake mode).

While the purpose of most security mechanisms used in the protocol seemed clear, actually performing the decomposition, i.e. finding suitable boundaries for “cutting” the protocol and specifying the assumed and provided guarantees at each layer, proved tedious and sometimes impossible; our approach was limited by many of TLS’ design choices. The protocol uses many techniques heuristically: the cryptographic keys intended to protect the payload are also used for the confirmation (here called “finished”) messages in TLS 1.2, the pseudo-random function is keyed with a key that is not a uniformly random bit string, and various protocol messages are unnecessarily included in the confirmation messages, to name a few.

These unfortunate design choices severely complicated our analysis, at multiple levels. The most straightforward example is that modules which (in the protocol description) consist of several steps cannot be analyzed in isolation. This is the case for the construction of the master secret key resource \mathbf{MSK} in TLS 1.2, which is constructed by a sub-protocol corresponding to the computation of both the premaster secret (PMS), and then of the master secret key (MSK)—in the terminology of [DR08]. This construction step cannot be modularized any further because the client (an untrusted entity, not authenticated in the case of unilateral authentication) chooses its input to the secret key after the server; intuitively, an attacker can at this point replace the client’s message by one that correlates the computed PMS with one in a different session, because the employed schemes are potentially malleable. As suggested in a previous version of this work, this issue is resolved in the current draft of TLS 1.3 [DR15]

by making the client choose its input to the key exchange first—a protocol of this type is also analyzed in [CMT13]). Alternatively, it would have been possible to use non-malleable (like CCA-secure) primitives for the key exchange.

Another complication arises from the fact that protocol sessions become distinguishable only at a late stage of the protocol, resulting in our resources providing for a multi-session scenario. Our resources for the augmented network resource **SNET** and the master secret key resource **MSK** have to provide for many sessions being run for pairs of nonces. This complication could be relieved by cryptographically binding a unique property of the client, like the nonce, to the messages exchanged in the protocol in earlier phases of the protocol. This would guarantee that “higher layers” of the protocol would not have to be analyzed in a multi-session scenario. Unfortunately, this criticism still applies to the current draft of TLS 1.3.

The fact that the “finished” messages use data from lower layers of the protocol (basically the entire session transcript of the lower layers) is reflected in the auxiliary information that is being passed through the layers, as well as in the numerous output values for each resource. Nevertheless, using the RSA-PKCS ciphertext as an input in the “finished” messages ensures that, in a monolithic analysis of the protocol, the OW-PCA assumption can be used (instead of NR-PCA which we need to assume here). In other words, this provides evidence that the lower level of modularity in the proof of [KPW13] for CCA security of the TLS key extraction mechanism is hard to avoid unless one is willing to make additional assumptions. Since our goal was to preserve modularity, we proved the security of the MSK-generation resource under the stronger NR-PCA assumption, introduced by [BFK⁺13b]. This issue is described more in detail in Section 3.3.3. A solution for this seems actually using an IND-CCA-secure encryption scheme for the key generation, rather than RSA-PKCS. Generally, values that are already authenticated (like the certificates or the signed messages) should not be included in the finished message at all; this complicates the analysis but does not contribute to the security of the protocol. The parameters used in the “finished” message generation and processing should be minimal so as to minimize the information passed from higher protocol layers. Furthermore, the well-known issue of using the encryption and MAC keys in the confirmation message prohibits to view the “finished messages” as part of the handshake protocol; rather, we view them as a first step in the record layer protocol. The first criticism, namely that several messages are included in the finished message unnecessarily, still applies to the current draft of TLS 1.3. The latter criticism, namely that the messages are protected with the same keys as the actual payload messages, does not apply anymore.

Our work also indicates that the sub-protocols cannot be fully decorrelated as long as TCP is used for message transmissions. As TCP fragments are processed in order-of-arrival, each of the resources has to permit further transmissions via two “insecure channels.” Finally, another questionable design choice of TLS 1.2 is the use of Authenticate-then-Encrypt, which is susceptible to timing- and other type of attacks [DP10]. This criticism has been addressed in TLS 1.3 by the adoption of AEAD ciphers. Still, as suggested by Badertscher et al. [BMM⁺15], by changing the way the AEAD cipher is used, the protocol efficiency can be improved.

Although various of the above mentioned obstacles could have been bypassed by modifying the scheme, we kept our analysis close to the realistic TLS protocol. The result is a more complicated analysis; however, the artifices are immanent to the protocol, not to our technique.

References

- [ABP⁺13] Nadhem J. AlFardan, Daniel J. Bernstein, Kenneth G. Paterson, Bertram Poettering, and Jacob C. N. Schuldt. On the security of RC4 in TLS and WPA. In *USENIX Security Symposium*, 2013.

- [AP12] Nadhem J. AlFardan and Kenneth G. Paterson. Plaintext-recovery attacks against datagram TLS. In *Network and Distributed System Security Symposium (NDSS'12)*, 2012.
- [AP13] Nadhem J. AlFardan and Kenneth G. Paterson. Lucky thirteen: Breaking the TLS and DTLS record protocols. In *IEEE Symposium on Security and Privacy (SP'13)*, 2013.
- [Bar04] Gregory V. Bard. Vulnerability of SSL to chosen-plaintext attack. Cryptology ePrint Archive: Report 2004/111, May 2004.
- [BFCZ12] Karthikeyan Bhargavan, Cédric Fournet, Ricardo Corin, and Eugen Zălinescu. Verified cryptographic implementations for TLS. In *ACM Transactions on Information and System Security (TISSEC'12)*, volume 15(1): 3, 2012.
- [BFK⁺13a] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, and Pierre-Yves Strub. Implementing TLS with verified cryptographic security. In *IEEE Symposium on Security and Privacy (SP'2013)*, pages 445–469, 2013.
- [BFK⁺13b] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Santiago Zanella-Béguelin. Proving the TLS handshake secure (as it is). Technical report, 2013. Available at http://www.mitls.org/downloads/Proving_the_TLS_Handshake.pdf.
- [BFS⁺13] Christina Brzuska, Marc Fischlin, Nigel Smart, Bogdan Warinschi, and Steve Williams. Less is more: Relaxed yet composable security notions for key exchange. *International Journal of Information Security*, 12(4):267–297, 2013.
- [BMM⁺15] Christian Badertscher, Christian Matt, Ueli Maurer, Phil Rogaway, and Björn Tackmann. Augmented secure channels as the goal of the TLS record layer. Manuscript, April 2015.
- [BPB12] Gilles Barthe, David Pointcheval, and Santiago Zanella Béguelin. Verified security of redundancy-free encryption from Rabin and RSA. In *ACM Conference on Computer and Communications Security*, pages 724–735, 2012.
- [BR93] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology — CRYPTO 1993*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. IACR, Springer, 1993.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145. IEEE, 2001. Extended version in [Can05].
- [Can05] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, December 2005.
- [CK01] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In *Proceedings of EUROCRYPT'01*, pages 453–474, 2001.
- [CK02] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In *Proceedings of EUROCRYPT'02*, pages 337–351, 2002.

- [CMT13] Sandro Coretti, Ueli Maurer, and Björn Tackmann. Constructing confidential channels from authenticated channels — Public-key encryption revisited. In *Advances in Cryptology — ASIACRYPT 2013*, Lecture Notes in Computer Science, Berlin Heidelberg, 2013. IACR, Springer.
- [DP10] Jean-Paul Degabriele and Kenneth G. Paterson. On the (in)security of ipsec in mac-then-encrypt. In *Proceedings of ACM Computer and Communications Security (ACM CCS) 2010*, 2010.
- [DR08] Tim Dierks and Eric Rescorla. The transport layer security (TLS) protocol version 1.2. RFC 5246, August 2008.
- [DR15] Tim Dierks and Eric Rescorla. The transport layer security (TLS) protocol version 1.3. RFC 5246 bis, April 2015.
- [FHM⁺12] Sascha Fahl, Marian Harbach, Thomas Muders, Lars Baumgärtner, Bernd Freisleben, and Matthew Smith. Why Eve and Mallory love Android: An analysis of Android SSL (in)security. In *Proceedings of the ACM Conference on Computer and Communications Security (ACM CCS’12)*, pages 50–61, 2012.
- [GIJ⁺12] Martin Georgiev, Subodh Iyengar, Suman Jana, Rishita Anubhai, Dan Boneh, and Vitaly Shmatikov. The most dangerous code in the world: Validating SSL certificates in non-browser software. In *Proceedings of the ACM Conference on Computer and Communications Security (ACM CCS’12)*, pages 38–49, 2012.
- [GKS13] Florian Giesen, Florian Kohlar, and Douglas Stebila. On the security of TLS renegotiation. Cryptology ePrint Archive Report 2012/630, 2013.
- [GMP⁺08] Sebastian Gajek, Mark Manulis, Olivier Pereira, Ahmad-Reza Sadeghi, and Jörg Schwenk. Universally composable security analysis of TLS. In *Proceedings of ProvSec 2008*, pages 313–327, 2008.
- [HFPS99] Russell Housley, Warwick Ford, Tim Polk, and David Solo. Internet X.509 public key infrastructure. RFC 2459, January 1999. Internet standard, <http://www.ietf.org/rfc/rfc2459.txt>.
- [Hic95] Kipp E. B. Hickman. The SSL protocol. Internet draft, February 1995.
- [HKR15] Viet Tung Hoang, Ted Krovetz, and Phillip Rogaway. Robust authenticated-encryption: Aez and the problem that it solves. In *Advances in Cryptology - EUROCRYPT 2015*, 2015. To appear.
- [HPFS02] Russell Housley, Tim Polk, Warwick Ford, and David Solo. Internet X.509 public key infrastructure; certificate and certificate revocation list (CRL profile). RFC 3280, April 2002. Internet standard, <http://www.ietf.org/rfc/rfc3280.txt>.
- [HSD⁺05] Changhua He, Mukund Sundararajan, Anupam Datta, Ante Derek, and John C. Mitchell. A modular correctness proof of IEEE 802.11i and TLS. In *Proceedings of the ACM Conference on Computer and Communications Security (ACM CCS’05)*, pages 2–15, 2005.
- [JK02] Jakob Jonsson and Burton S. Kaliski Jr. On the security of RSA encryption in TLS. In *Proceedings of CRYPTO 2002*, pages 127–142, 2002.

- [JKSS12] Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In *Proceedings of CRYPTO 2012*, pages 273–293, 2012.
- [Kal98] Burt Kaliski. PKCS #7: Cryptographic message syntax. RFC 2315, March 1998. Version 1.5.
- [KMO⁺13] Markulf Kohlweiss, Ueli Maurer, Cristina Onete, Björn Tackmann, and Daniele Venturi. Anonymity-preserving public-key encryption: A constructive approach. In E. De Cristofaro and M. Wright, editors, *PETS 2013*, volume 7981 of *LNCS*, pages 19–39, Heidelberg, 2013. Springer.
- [KPW13] Hugo Krawczyk, Kenneth Paterson, and Hoeteck Wee. On the security of the TLS protocol: A systematic analysis. In *Proceedings of CRYPTO 2013*, pages 429–448, 2013.
- [Kra01] Hugo Krawczyk. The order of encryption and authentication for protecting communication (or: How secure is SSL?). In *Proceedings of CRYPTO 2001*, pages 310–331, 2001.
- [KSS13] Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DH and TLS-RSA in the standard model, 2013.
- [KT11] Ralf Küsters and Max Tuengerthal. Composition theorems without pre-established session identifiers. Cryptology ePrint Archive, Report 2011/406, 2011.
- [Mau02] Ueli Maurer. Indistinguishability of random systems. In Lars R. Knudsen, editor, *Advances in Cryptology — EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 110–132. IACR, Springer-Verlag, 2002.
- [Mau11] Ueli Maurer. Constructive cryptography: A new paradigm for security definitions and proofs. In Sebastian Mödersheim and Catuscia Palamidessi, editors, *TOSCA 2011—Theory of Security and Applications*, Lecture Notes in Computer Science. Springer-Verlag, 2011.
- [Mau13] Ueli Maurer. Conditional equivalence of random systems and indistinguishability proofs. In *2013 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 3150–3154, July 2013.
- [MR11] Ueli Maurer and Renato Renner. Abstract cryptography. In *Innovations in Computer Science*. Tsinghua University Press, 2011.
- [MSW08] Paul Morrissey, Nigel Smart, and Bogdan Warinschi. A modular security analysis of the TLS handshake protocol. In *Proceedings of ASIACRYPT 2008*, pages 55–73, 2008.
- [MT10] Ueli Maurer and Björn Tackmann. On the soundness of Authenticate-then-Encrypt: Formalizing the malleability of symmetric encryption. In *ACM Conference on Computer and Communications Security*. ACM, 2010.
- [MTC13] Ueli Maurer, Björn Tackmann, and Sandro Coretti. Key exchange with unilateral authentication: Composable security definition and modular protocol design. Cryptology ePrint Archive, Report 2013/555, 2013. <http://eprint.iacr.org/>.

- [OP01] Tatsuaki Okamoto and David Pointcheval. The gap problems: A new class of problems for the security of cryptographic schemes. In *PKC*, volume 1992 of *LNCS*, pages 104–118. Springer, 2001.
- [Pau99] Lawrence C. Paulson. Inductive analysis of the internet protocol TLS. In *ACM Transactions on Information and System Security (TISSEC)*, pages 332–351, 1999.
- [PRS11] Kenneth G. Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology — ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 372–389. IACR, Springer-Verlag, 2011. To appear.
- [SCF⁺11] Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. Secure distributed programming with value-dependent types. In *ICFP*, pages 266–278, 2011.
- [Sho99] Victor Shoup. On formal models for secure key exchange. Research Report RZ 3120, IBM, April 1999.
- [WS96] David Wagner and Bruce Schneier. Analysis of the SSL 3.0 protocol. In *USENIX Workshop on Electronic Commerce*, pages 29–40, 1996.

A More Details on TLS

A.1 X.509 Certificates

In TLS, the server’s public key is always certified; optionally, the client may also be required to certify his public key. The certificates used by TLS are X.509 v3 certificates [HPFS02], in practice a chain of certificates starting from the user and ending at a valid certification authority. In practice, each X509 v3 certificate consists of a sequence of three required fields: the `TBSCertificate`, the `AlgorithmIdentifier`, and the `BIT STRING` fields (the latter being a signature). We give more details about the concrete certificate structure below.

In this paper, we abstract the certification process to (user access to) a resource $\text{PKI}_{\mathfrak{F}}$, which takes as input the values that need to be certified, and outputs a certificate, i.e. a bit-string which can be verified. The resource will choose a function $f \in \mathfrak{F}$ and output a certificate of the form $(x, f(x))$, where x is the input that needs to be certified. The function f can be seen as an abstraction of the precise signature generation algorithm. Note that in the X509 certificates the values that are certified are encoded in DER encoding, in particular ensuring that the encoding yields a *unique* value. The encoding consists of a *type*, the *length* of the payload, and the *payload* itself, i.e. the input. We note that this encoding is implicitly assumed in our PKI resource, i.e. the input value x is assumed to correspond to a single certificate. Furthermore, the `tbsCertificate` field contains a unique identifier which bounds each input to a single certificate (per resource), which means, the certificate authority can check whether it has issued the certificate or not.

Finally, we note that in our assessment, we simplify the output of the verification of a certificate to a single bit, i.e. the certificate is valid or invalid. This is a simplification since in fact the output usually yields more information than this: there is a difference between e.g. a valid, but revoked certificate, a certificate which has expired, and an invalid certificate.

We proceed to describe the concrete structure of the X.509 v3 certificates with three main fields: `TBSCertificate`, `AlgorithmIdentifier`, and `BIT STRING`. The field `TBSCertificate` consists of:

- **Version.** The version of the certificate, i.e. 1, 2, or 3. If the version is 1, the value is omitted. Already in version 2, the certificates had a unique identifier.
- **Serial number.** This is a unique serial number, which is a positive integer, 160 bits long.
- **Signature.** This field contains the signing algorithm, which must be the same as the one specified in the `AlgorithmIdentifier` field; it may also contain additional parameters.
- **Issuer.** The name of the issuing identity, specified as a sequence of attributes such as country, organization, state or province name, common name, serial number, etc.
- **Validity.** A sequence of two dates, a start and an end date for the validity.
- **Subject.** This field contains the (unique per each issuing CA) identifier of the owner of the public key to be certified.
- **Subject public key info.** This sub-field contains the public key and the algorithm with which the key should be used (e.g. RSA or Diffie-Hellman).
- **Unique identifiers.** These are the subject and issuer unique identifiers, which allow the certification to handle repetitions of either the subject or the issuer fields.
- **Extensions.** The certificates generated in version 3 feature this field, which consists of a sequence of one or more certificate extensions. Examples of such extensions are: subject key identifier (which allows to identify certificates containing a certain public key), key usage (which restricts the use of the key for particular purposes, such as digital signatures, key agreement, etc.), subject alternative name (enabling additional identities to be bound to the subject of the certificate), etc.

The `AlgorithmIdentifier` value contains the algorithm used to generate the signature. The third field, BIT STRING is the signature over the DER encoding of the `TBSCertificate` field.

A.2 RSA PKCS#7

A public-key encryption (PKE) scheme with message space \mathcal{M} is typically described as three algorithms $PKE = (gen, enc, dec)$. The key-generation algorithm gen outputs a key pair (pk, sk) , the (probabilistic) encryption algorithm enc takes a message $m \in \mathcal{M}$ and a public key pk and outputs a ciphertext $c = enc(m; pk)$, and the decryption algorithm takes a ciphertext c and a secret key sk and outputs a plaintext $m = dec(c; sk)$. It is possible that the output of the decryption algorithm is the special symbol \perp ; this indicates that the ciphertext c is invalid.

Below, we describe the RSA algorithms following the standard PKCS#7 [Kal98]. Let $\lambda_0 = \Theta(\lambda)$, $\lambda_1 = \Theta(\lambda)$ with $\lambda_0 \leq \lambda_1 - 88$. It is also assumed that λ_1 is a multiple of 8. Strictly speaking, PKCS#7 does not specify a key-generation algorithm gen , but assumes that some RSA key pair is already available. Hence, all our security statements with respect to RSA-based cipher suites are implicitly parametrized by the actual algorithm gen that was used to generate the server's key pair. Consider the following triple of algorithms $RSA = (gen, enc, dec)$.

- $gen(\lambda)$: Upon input the security parameter λ , output $(pk, sk) = ((M, e), d)$ such that $ed \equiv 1 \pmod{\phi(M)}$ and the modulus M has λ_1 bits.
- $enc(m; pk)$: Upon input a λ_0 -bit message m , pick a random padding $P \in \{0, 1\}^{\lambda_1 - \lambda_0 - 24}$ (such that none of the bytes of P equal '00' in hexadecimal notation), define $x = 00||02||P||00||m$, and output $c = x^e \pmod{M}$.
- $dec(c; sk)$: Upon input c , attempt to parse $c^d \pmod{M}$ as a sequence of bytes of the form $00||02||P||00||m$ such that P contains no zero bytes and m has exactly λ_0 bits. If the attempt is successful output m ; otherwise output a special symbol \perp .

When used within TLS, λ_0 is fixed to 384 (yielding a pre-master secret of 48 bytes), while λ_1 is typically 1024 or 2048.

OW-PCA and NR-PCA. The RSA-PKCS-based version of the TLS protocol uses the above defined PKE scheme. Earlier works analyzing the security of TLS-RSA relied on the assumption that RSA-PKCS satisfies a special property called *One-Wayness under Plaintext Checking Attacks* (OW-PCA). This notion is reviewed in Appendix C.1. Not much is known on the validity of this assumption for RSA PKCS#7. The only relevant paper trying to justify it is [JK02], via a reduction to an RSA-like assumption (a.k.a. partial-domain RSA with decision oracle).²³

Similar to [BFK⁺13b], our security proof for RSA-TLS relies on a stronger assumption called *Non-Randomizability under Plaintext Checking Attacks*. As in the case of OW-PCA not much is known on the validity of this assumption for RSA PKCS#7. See also Appendix C.1 for a discussion.

A.3 Key Expansion

The key expansion procedure used in TLS relies on a pseudo-random function (based on HMAC). This PRF with the SHA-256 hash function is used for all cipher suites defined in TLS 1.2 (see Section 4). To expand the key, the following function is defined taking as input a secret, a seed, and an identifying label (and produces an output of arbitrary length):

$$\begin{aligned} P_{\text{hash}}(\text{secret}, \text{seed}) = & \text{HMAC}_{\text{hash}}(\text{secret}, A(1) + \text{seed}) + \text{HMAC}_{\text{hash}}(\text{secret}, A(2) + \text{seed}) \\ & + \text{HMAC}_{\text{hash}}(\text{secret}, A(3) + \text{seed}) + \dots \end{aligned}$$

where $+$ indicates concatenation. The function $A(\cdot)$ is defined as: $A(0) = \text{seed}$ and $A(i) = \text{HMAC}_{\text{hash}}(\text{secret}, A(i-1))$. Note that P_{hash} can be iterated as many times as necessary to produce the required quantity of data.

B Further Notation and Preliminaries

This section contains further definitions which have been deferred from the main body of the paper. The following Lemma, copied from [Mau13], states that if two games are equivalent, the probability of winning is the same.

Lemma 25 (Mau13, Lemma 1). *If $\mathbf{S} \stackrel{g}{=} \mathbf{T}$, then for any system \mathbf{D} and any q ,*

$$\Gamma_q^{\mathbf{D}}(\mathbf{S}) = \Gamma_q^{\mathbf{D}}(\mathbf{T}).$$

More importantly, the following lemma states that if two systems are equivalent as games, then the distinguishing advantage is upper bounded by the probability of winning the games. This lemma, which originates from [Mau02], is instrumental for many of our proofs.

Lemma 26 (Mau13, Lemma 2). *Let \mathcal{A} be a MBO. If $\mathbf{S}^{\mathcal{A}} \stackrel{g}{=} \mathbf{T}^{\mathcal{A}}$, then, for any distinguisher \mathbf{D} and any q ,*

$$\Delta_q^{\mathbf{D}}(\mathbf{S}, \mathbf{T}) \leq \Gamma_q^{\mathbf{D}}(\mathbf{S}^{\mathcal{A}}).$$

The following lemma is also needed by some of our proofs. The intuitive interpretation is as follows: For a tuple of games $\mathbf{G}_1, \dots, \mathbf{G}_n$ which have individually defined MBOs but are equivalent as games *with respect to the disjunction of their MBOs*, the sum of advantages of winning the individual games is at least as large as the advantage for provoking the disjunction.

²³It is not clear whether the result of [JK02] applies to RSA PKCS#7 with typical parameters as used in TLS. However we remark that, since its introduction, no weaknesses on the assumption have been reported either.

Lemma 27. Let $\mathbf{G}_1, \dots, \mathbf{G}_n$ be a family of random systems (i.e., each \mathbf{G}_i is described by a family of $\mathbf{p}_{Y_q|X_q}^{\mathbf{G}_i}$ for $q \geq 1$) and $\mathcal{A}^1, \dots, \mathcal{A}^n$ be a family of monotone binary outputs defined on these systems such that A_q^i is independent of Y_q . For $i = 1, \dots, n$, we write each monotone output as $\mathcal{A}^i = (A_1^i, A_2^i, \dots)$. Define $\mathcal{A} = \bigvee_{i=1}^n \mathcal{A}^i$, with $\mathcal{A} = (A_1, A_2, \dots)$ and $A_q = \bigvee_{i=1}^n A_q^i$ (i.e., A_q becomes 1 as soon as there exists a monotone output \mathcal{A}^i whose q^{th} component is 1). Assume that $\mathbf{G}_i^{\mathcal{A}} \stackrel{g}{=} \mathbf{G}_j^{\mathcal{A}}$ for all $1 \leq i, j \leq n$. Then, for all adversaries \mathbf{A} ,

$$\sum_{i=1}^n \Gamma_q^{\mathbf{A}}(\mathbf{G}_i^{\mathcal{A}^i}) \geq \Gamma_q^{\mathbf{A}}(\mathbf{G}_{i^*}^{\mathcal{A}}),$$

for any $1 \leq i^* \leq n$.

Proof. We define the monotone binary outputs $\tilde{\mathcal{A}}^1, \dots, \tilde{\mathcal{A}}^n$ as

$$\tilde{\mathcal{A}}_q^i = \tilde{\mathcal{A}}_{q-1}^i \vee \left(A_q^i \wedge \neg \left(\bigvee_{j \neq i} A_{q-1}^j \right) \right),$$

i.e., $\tilde{\mathcal{A}}^i$ formalizes that \mathcal{A}^i is (among) the *first* MBOs to become 1. Indeed, $\tilde{\mathcal{A}}_q^i = 1$ if either: (1) the MBO $\tilde{\mathcal{A}}^i$ had already turned 1, i.e., $\tilde{\mathcal{A}}_{q-1}^i = 1$; (2) the q^{th} value A_q^i is 1, but no previous value A_{q-1}^j of any other MBO \mathcal{A}^j is 1. Still, \mathcal{A} becomes 1 as soon as any one of the outputs A_q^i becomes 1 for some i, q . If this output becomes 1, there must be at least one $\tilde{\mathcal{A}}_q^i$ that became 1 first. Thus, \mathcal{A} becomes 1 if and only if at least one $\tilde{\mathcal{A}}_q^i$ becomes 1, yielding $\mathcal{A} = \bigvee_{i=1}^n \tilde{\mathcal{A}}^i$. Let now $1 \leq i^* \leq n$. It holds that: $\sum_{i=1}^n \Gamma_q^{\mathbf{A}}(\mathbf{G}_{i^*}^{\tilde{\mathcal{A}}^i}) \geq \Gamma_q^{\mathbf{A}}(\mathbf{G}_{i^*}^{\mathcal{A}})$, since \mathcal{A} is triggered if and only if it was triggered first in one of the $\tilde{\mathcal{A}}^i$ outputs. On the other hand, we also have $\Gamma_q^{\tilde{\mathcal{A}}^i}(\mathbf{G}_i) = \Gamma_q^{\tilde{\mathcal{A}}^i}(\mathbf{G}_j)$ since provoking $\tilde{\mathcal{A}}$ implies provoking \mathcal{A} while \mathbf{G}_i and \mathbf{G}_j are still behaving equivalently (since, if $\tilde{\mathcal{A}}$ was not triggered before, this means that \mathcal{A} does not hold). Hence, we obtain

$$\sum_{i=1}^n \Gamma_q^{\mathbf{A}}(\mathbf{G}_i^{\mathcal{A}^i}) \geq \sum_{i=1}^n \Gamma_q^{\mathbf{A}}(\mathbf{G}_i^{\tilde{\mathcal{A}}^i}) = \sum_{i=1}^n \Gamma_q^{\mathbf{A}}(\mathbf{G}_{i^*}^{\tilde{\mathcal{A}}^i}) \geq \Gamma_q^{\mathbf{A}}(\mathbf{G}_{i^*}^{\mathcal{A}}),$$

which concludes the proof. \square

B.1 Signature Schemes

A signature scheme is a triple of algorithms $SIG = (gen, sign, vrf)$. The *key-generation* algorithm gen takes no input²⁴ and outputs a pair (sk, vk) of a *signature key* sk and a *verification key* vk . The *signing* algorithm $sign$ takes as input a signature key sk and a message $m \in \mathcal{M}$ of some message space \mathcal{M} , and outputs a signature $s = sign(sk, m)$. The (often deterministic) *verification* algorithm vrf takes as input a verification key vk , a message m , and a signature s , and outputs a decision bit. A signature scheme is correct if for any key pair (sk, vk) generated by gen and for all $m \in \mathcal{M}$, $vrf(vk, m, sign(sk, m)) = 1$.

The common security requirement for a signature scheme $SIG = (gen, sign, vrf)$ is called *unforgeability* and is formalized in Section C.4.

using the following game \mathbf{G}^{SIG} :

1. Generate a key pair $(sk, vk) = gen()$ and output vk to the adversary.

²⁴For an asymptotic treatment, the algorithm takes as input the security parameter.

2. (Repeatedly) Given a message $m \in \mathcal{M}$ from the adversary, compute $s = \text{sign}(sk, m)$, store m in an internal buffer \mathcal{B} , and return s to the adversary.
3. Upon input a pair (m', s') with $m' \notin \mathcal{B}$ and $\text{verf}(vk, m', s') = 1$, output that the game is won.

For $\varepsilon \in [0, 1]$, a signature scheme is ε -secure with respect to a class \mathcal{D} of adversaries if $\Gamma^{\mathbf{A}}(\mathbf{G}^{\text{SIG}}) \leq \varepsilon$ for all $\mathbf{A} \in \mathcal{D}$.

C Game-based Definitions

This Appendix collects the relevant game-based definitions that are used in our analysis of TLS.

Game-based definitions specify a property of a cryptographic scheme based on an interaction between two (hypothetical) entities: the game (or challenger) and the adversary. During the interaction, the adversary may issue “oracle queries” to the challenger, the responses of which model what information may be leaked to the adversary. The adversary’s goal is specified by the game, and could be, e.g., forging a message or distinguishing encryptions of different messages. If this game cannot be won by any (efficient) adversary, then the scheme is secure against the considered type of attack. The formal definition of a game is given in Definition 6.

Bit guessing games Some games in the literature are *bit-guessing games*. These games can often be described by a pair of systems \mathbf{G}_0 and \mathbf{G}_1 , with the interpretation that in the beginning of the game, a bit $B \in \{0, 1\}$ is chosen uniformly at random. The adversary will then be given access to \mathbf{G}_B , and the goal is to guess the bit B . The adversary can win such a game with probability $\frac{1}{2}$ trivially by simply guessing the hidden bit. Hence, we measure the adversary’s success in terms of his *advantage*, that is, the (absolute) difference between \mathbf{A} ’s probability of winning \mathbf{G} and the success probability for these “trivial” strategies, formally $\Phi^{\mathbf{A}}(\mathbf{G}) = 2 \cdot |\Gamma^{\mathbf{A}}(\mathbf{G}) - \frac{1}{2}|$. Note also that $\Phi^{\mathbf{A}}(\mathbf{G}) = \Delta^{\mathbf{A}}(\mathbf{G}_0, \mathbf{G}_1)$.

C.1 OW-PCA & NR-PCA

We review the notion of one-wayness against plaintext checking attacks [JK02]. Let $PKE = (\text{gen}, \text{enc}, \text{dec})$ be a public key encryption scheme with message space \mathcal{M} .

Consider the game of Figure 7.

Init()	ChGen()	PCA(m, c)	GameOutput(m')
$(pk, sk) \leftarrow \text{gen}()$	if $\text{Chal} \neq \emptyset$	if $(m = \perp) \vee (\text{Chal} = \emptyset)$	if $\text{Chal} = \emptyset$
$\text{Chal} \leftarrow \emptyset$	return \perp	return \perp	return \perp
$W \leftarrow 0$	else	else if $m = \text{dec}(c; sk)$	else
return pk	$m^* \leftarrow \mathcal{M}$	return 1	return $\text{Output} = (m' = m^*)$
end.	$\text{Chal} \leftarrow \text{enc}(m^*; pk)$	else	end.
	return Chal	return 0	
	end.	end.	

Figure 7: The OW-PCA security game, $\mathbf{G}^{\text{ow-pca}}$

Definition 28. The encryption scheme $PKE = (\text{gen}, \text{enc}, \text{dec})$ is ε -OW-PCA with respect to a class \mathcal{D} of adversaries if for every $\mathbf{A} \in \mathcal{D}$ it holds that $\Gamma^{\mathbf{A}}(\mathbf{G}^{\text{ow-pca}}) \leq \varepsilon$.

Definition 28 intuitively says that it is hard to “invert” a ciphertext, even given access to a plaintext checking oracle (i.e., an oracle allowing to check if a guess for the plaintext

corresponding to some ciphertext is correct). As discussed in Section 3.3.3, our proof for TLS-RSA also relies on the assumption that it is hard to re-randomize a ciphertext, even given access to a plaintext checking oracle. The latter notion, also known as non-randomizability against plaintext checking attacks (NR-PCA), was recently introduced in [BFK⁺13b]; the corresponding game is depicted in Figure 8.

Init()	ChGen()	PCA(m, c)	GameOutput(c')
$(pk, sk) \leftarrow \text{gen}()$ $\text{Chal} \leftarrow \emptyset$ $W \leftarrow 0$ return pk end.	if $\text{Chal} \neq \emptyset$ return \perp else $m^* \leftarrow \mathcal{M}$ $\text{Chal} \leftarrow \text{enc}(m^*; pk)$ return Chal end.	if $(m = \perp) \vee (\text{Chal} = \emptyset)$ return \perp else if $m = \text{dec}(c; sk)$ return 1 else return 0 end.	if $\text{Chal} = \emptyset$ return \perp else return $\text{Output} = (c' \neq c \wedge \text{dec}(c'; sk) = m^*)$ end.

Figure 8: The NR-PCA security game, $\mathbf{G}^{\text{nr-pca}}$

Definition 29. The encryption scheme $PKE = (\text{gen}, \text{enc}, \text{dec})$ is ε -NR-PCA with respect to a class \mathcal{D} of adversaries if for every $\mathbf{A} \in \mathcal{D}$ it holds that $\Gamma^{\mathbf{A}}(\mathbf{G}^{\text{nr-pca}}) \leq \varepsilon$.

Note that NR-PCA implies OW-PCA whenever the encryption algorithm is randomized, because if we can invert the challenge ciphertext we can also re-randomize it. The other direction might not be true, as there might be easier ways to re-randomize a ciphertext than by inverting it. [BFK⁺13b] conjectured that the NR-PCA assumption follows from the common-input extractability assumption of [BPB12] and OW-PCA.

In TLS the message space is of the form $\mathcal{M} = \text{version_number} \times \{0, 1\}^{368}$. Note that decryption may result in a message that is outside of the message space and that such invalid plaintexts can still be checked using the **PCA** oracle.

C.2 Gap Diffie-Hellman

For some security statements, we use the gap Diffie-Hellman assumption, originally proposed by Okamoto and Pointcheval [OP01], which essentially states that it is hard to compute g^{xy} given (g, g^x, g^y) , even given access to a DDH verification oracle (\cdot, \cdot, \cdot) . We formalize this problem parametrized by a distribution \mathcal{G} over groups of finite order $|\mathbb{G}| = q$, together with a (publicly-known) generator g . The game is specified in Figure 9.

Init(\mathcal{G})	DDH(g^a, g^b, C)	GameOutput(Z)
draw $(\mathbb{G}, g) \leftarrow \mathcal{G}$ draw $x, y \leftarrow \{1, \dots, \mathbb{G} \}$ return $(\mathbb{G}, g, g^x, g^y)$ set $W \leftarrow 0$ end.	return $(g^{ab} = C)$ end.	$W \leftarrow (Z = g^{xy})$ end.

Figure 9: The GapDH security game, $\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}$

Definition 30. A group distribution \mathcal{G} ε -satisfies the Gap Diffie-Hellman assumption with respect to the class \mathcal{D} of adversaries and with error ε if for all $\mathbf{A} \in \mathcal{D}$: $\Gamma^{\mathbf{A}}(\mathbf{G}_{\mathcal{G}}^{\text{GapDH}}) \leq \varepsilon$.

C.3 Collision Resistance

For the key expansion step in Section 4, the client and server use a hash function H . In order to prevent the adversary from changing some of the protocol messages without modifying the finished messages ξ_C, ξ_S , this hash function needs to be collision resistant.

$$\frac{\mathbf{GameOutput}(x, x')}{W \leftarrow (x = x')}$$

end.

Figure 10: The collision resistance game, \mathbf{G}^{CR}

Definition 31. A hash function H is ε -collision-resistant for a class \mathcal{D} of adversaries, if for all $\mathbf{A} \in \mathcal{D}$ it holds: $\Gamma^{\mathbf{A}}(\mathbf{G}^{\text{CR}}) \leq \varepsilon$.

C.4 Unforgeability under Chosen-Message Attacks

The security condition for a signature scheme as defined in Section B.1 is a tuple of algorithms $SIG = (gen, sign, vrf)$. The standard security requirement for a signature scheme is *unforgeability under chosen-message attacks* (*UF-CMA*) and formalizes that no attacker may be able to forge a signature, even if he is given access to a “signature oracle” that returns signatures for arbitrary messages. (Of course, signatures returned by the oracle are not eligible for winning the game.)

Init()	Sign(m)	Forge(m, s)
$(sk, vk) \leftarrow gen()$	$s \leftarrow sign(m; sk)$	if $(m \notin \mathcal{B}) \wedge vrf(m, s; vk)$
$\mathcal{B} \leftarrow \emptyset$	$\mathcal{B} \leftarrow \mathcal{B} \cup \{m\}$	$W \leftarrow 1$
$W \leftarrow 0$	return s	end.
return vk	end.	
end.		

Figure 11: The unforgeability game for the scheme SIG , $\mathbf{G}^{\text{uf-cma}}$.

Definition 32. A signature scheme is ε -existentially unforgeable under chosen message attacks for a class \mathcal{D} of adversaries, if for all $\mathbf{A} \in \mathcal{D}$ it holds: $\Gamma^{\mathbf{A}}(\mathbf{G}^{\text{uf-cma}}) \leq \varepsilon$.

C.5 Authenticated Encryption with Associated Data

We define the security game for AEAD-schemes using the all-in-one formulation of Hoang et al. [HKR15]. A scheme is considered secure if all valid and efficient adversaries have poor advantage according to the following definition.

We describe two systems, a “real” one and an “ideal” one. In the “real” system, encryption (resp. decryption) queries to the scheme are answered by encrypting (resp. decrypting) the given message (resp. ciphertext) with respect to the given nonce and associated data. In the “ideal” system, encryption queries are answered with a uniformly random string of appropriate length (it is required that this only depends on the length—not the value—of the plaintext); decryption queries are either answered with the corresponding plaintext (if they correspond to outputs of a previous encryption query) or by a special “invalid” symbol. We refer to those systems as $\mathbf{G}_0^{\text{aead}}$ and $\mathbf{G}_1^{\text{aead}}$, respectively.

Init()	Enc (n, a, m)	Dec (n, a, c)
$k \leftarrow \mathcal{K}$ $\mathcal{B} \leftarrow \emptyset$ end.	$c \leftarrow \text{enc}(k, n, a, m)$ $\mathcal{B} \leftarrow \mathcal{B} \cup \{c\}$ return c end.	if $c \notin \mathcal{B}$ $m \leftarrow \text{dec}(k, n, a, c)$ return m end.
Init()	Enc (n, a, m)	Dec (n, a, c)
$k \leftarrow \mathcal{K}$ $\mathcal{B} \leftarrow \emptyset$ end.	$c \leftarrow \{0, 1\}^{\ell(m)}$ $\mathcal{B} \leftarrow \mathcal{B} \cup \{c\}$ return c end.	if $c \notin \mathcal{B}$ $m \leftarrow \perp$ return m end.

Figure 12: The security game for authenticated encryption, above $\mathbf{G}_0^{\text{aead}}$, below $\mathbf{G}_1^{\text{aead}}$.

Definition 33. An AEAD scheme is ε -secure for a class \mathcal{D} of adversaries if for all $\mathbf{A} \in \mathcal{D}$, it holds that $\Delta^{\mathbf{A}}(\mathbf{G}_0^{\text{aead}}, \mathbf{G}_1^{\text{aead}}) \leq \varepsilon$.

Resource	Description	Constructed from	Sub-protocol	Security Loss
Basic Resources:				
\rightarrow	one-way, insecure channel; can read, change, and inject messages at E -interface.	assumed resource		0
$\text{PKI}_{\mathfrak{F}}$	Public-Key Infrastructure; provides certification of PKs.	assumed resource		0
RO_{384}	Random Oracle; outputs consistent randomness.	assumed resource		0
NET	the insecure point-to-point network; parallel composition of \rightarrow and \leftarrow channels; behave as for \rightarrow at E -interface.	assumed resource, $\llbracket [-\rightarrow, \leftarrow -] \rrbracket^{(C, S/C)}$ for $C \in \mathcal{A}_{\text{TCP}}$		0
NAME_{ρ}	unique name resource; associates each client interface $C \in \mathcal{C}$ with a unique nonce η , as ρ indicates.		rnd	$\binom{ \mathcal{C} }{2} \cdot 2^{-224}$
Intermediate TLS Resources				
$\text{SNET}_{\text{N}, \rho, n}$	insecure network with nonce exchange; associates each client interface $C \in \mathcal{C}$ with a unique nonce η , as ρ indicates.	$\text{NAME}_{\rho}, \text{NET}$	hec, hes	$\binom{n}{2} \cdot 2^{-224}$
$\xrightarrow{\bullet} \text{N}_{\rho, \mathfrak{F}, \text{SIG}, n}$	authenticated transmission network; network with one-sided authentication of the group parameters used in TLS-DHE.	$\text{SNET}_{\text{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}$	vrf, sgn	$\Gamma^{\text{DC}}(\mathbf{G}^{\text{uf-cma}})$
$\text{MSK}_{\text{N}, \rho, \text{AUX}, n}$	master secret resource; allows parties to obtain the master secret. For each session, MS can either be injected or honestly generated.	TLS-DH: $\text{RO}_{384}, \text{SNET}_{\text{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}$ TLS-DHE: $\text{RO}_{384}, \xrightarrow{\bullet} \text{N}_{\rho, \mathfrak{F}, \text{SIG}, n}$ TLS-RSA: $\text{RO}_{384}, \text{SNET}_{\text{N}, \rho, n}, \text{PKI}_{\mathfrak{F}}$	dhc, dhs _G . dhec, dhes _G rsac, rsas	$\Gamma^{\text{DC}}(\mathbf{G}_G^{\text{GapDH}})$ $n \mathcal{C} \cdot \Gamma^{\text{DC}}(\mathbf{G}_G^{\text{GapDH}}) + \Gamma^{\text{DC}'}(\mathbf{G}^{\text{uf-cma}})$ $n \mathcal{C} \cdot \Gamma_q^{\text{DC}'}(\mathbf{G}^{\text{nr-pca}}) + \frac{q}{2^{368}}$
Final TLS Resources:				
$\xleftarrow{\text{KSP},*} \bullet$	unilaterally authenticated key; can inject keys or allow them to be honestly distributed.	$\text{MSK}_{\text{N}, \rho, \text{AUX}, n}$	expc, exps	$ \mathcal{C} \cdot \Gamma^{\text{DC}}(\mathbf{G}^{\text{PRF}}) + \Gamma^{\text{DC}'}(\mathbf{G}^{\text{CR}})$
$\xleftarrow{*} \xrightarrow{\bullet}$	Unilaterally secure 2-party communication. For each session, adversary can either interfere or not.	$\xleftarrow{\text{KSP},*} \bullet$	stream cipher CBC	$2 \cdot \Delta^{\text{DC}}(\text{stream } \mathbf{U}_k, \mathbf{U}^*) + 2 \cdot \Gamma^{\text{DC}'}(\mathbf{G}^{\text{suf-cma}})$ $2 \cdot \Delta^{\text{DC}}(\text{bc } \mathbf{U}_k, \mathbf{P}_{\ell}) + 2 \cdot \Gamma^{\text{DC}'}(\mathbf{G}^{\text{suf-cma}}) + \frac{(q\ell)^2}{2^{\ell-1}}$

Figure 13: The resources used in this work